

ClusterPack



Index of Tutorial Sections

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

[Administrators Guide](#)

- [1.0 ClusterPack Install QuickStart](#)
- [1.1 ClusterPack General Overview](#)
- [1.2 Comprehensive Install Instructions](#)
- [1.3 Installation and Configuration of Optional Components](#)
- [1.4 Software Upgrades and Reinstalls](#)
- [1.5 Golden Image Tasks](#)
- [1.6 System Maintenance Tasks](#)
- [1.7 System Monitoring Tasks](#)
- [1.8 Workload Management Tasks](#)
- [1.9 System Troubleshooting Tasks](#)

[Users Guide](#)

- [2.1 Job Management Tasks](#)
- [2.2 File Transfer Tasks](#)
- [2.3 Miscellaneous Tasks](#)

[Tool Overview](#)

- [3.1 Cluster Management Utility Zone Overview](#)
- [3.2 Service ControlManager \(SCM\) Overview](#)
- [3.3 System Inventory Manager \(SIM\) Overview](#)
- [3.4 Application ReStart \(AppRS\) Overview](#)
- [3.5 Cluster Management Utility \(CMU\) Overview](#)
- [3.6 NAT/IPFilter Overview](#)
- [3.7 Platform Computing Clusterware Pro V5.1 Overview](#)
- [3.8 Management Processor \(MP\) Card Interface Overview](#)

[Related Documents](#)

- [4.1 Related Documents](#)

[Dictionary of Cluster Terms](#)

[Back to Top](#)

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

Copyright 1994-2004 hewlett-packard company

ClusterPack



ClusterPack Install QuickStart

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

[1.0.1 How Can I Get My HP-UX Cluster Running?](#)

[Step Q1 Fill Out the ClusterPack Installation Worksheet](#)

[Step Q2 Install Prerequisites](#)

[Step Q3 Allocate File System Space](#)

[Step Q4 Obtain a License File](#)

[Step Q5 Prepare Hardware Access](#)

[Step Q6 Power Up the Management Server](#)

[Step Q7 Configure the ProCurve Switch](#)

[Step Q8 Copy the License Files to the Management Server](#)

[Step Q9 Install ClusterPack on the Management Server](#)

[Step Q10 Run manager_config on the Management Server](#)

[Step Q11 Run mp_register on the Management Server](#)

[Step Q12 Power up the Compute Nodes](#)

[Step Q13 Run compute_config on the Management Server](#)

[Step Q14 Run finalize_config on the Management Server](#)

1.0.1 How Can I Get My HP-UX Cluster Running?

If you have installed ClusterPack before, follow the instructions in this section as a quick reminder. You can refer to the detailed instructions for any given step via the links provided.

If you have not installed ClusterPack before, use the Comprehensive Install Instructions section instead. This QuickStart Guide does NOT cover the use of Golden Images. If you wish to use Golden Images, use the Comprehensive Install Instructions section instead.

Note:

If you do the installation steps out of order or omit steps, your installation will leave your systems in an unknown and non-deterministic state.

[Back to Top](#)

Step Q1 Fill Out the ClusterPack Installation Worksheet

Print out this form and fill out all information for each node in your cluster.

[Installation Worksheet \(pdf\)](#)

Note:

You will not be able to complete the following steps if you have not collected all of this information.

For more information, see the Comprehensive Instructions for this step.

References:

- [Step 1 Fill Out the ClusterPack Installation Worksheet](#)

[Back to Top](#)

Step Q2 Install Prerequisites

Install the following software on the Management Server.

- HP-UX 11i Ignite-UX
- HP-UX 11i V2.0 TCOE

Install the following software on each Compute Node.

- HP-UX 11i Ignite-UX
- HP-UX 11i V2.0 TCOE

Allow the default choices to install.

ClusterPack requires a homogeneous operating system environment. That is, all Compute Nodes and the Management Server must have the same release of HP-UX installed as well as the same operating environment.

The Management Server requires a minimum of two LAN connections. One connection must be configured prior to installing ClusterPack.

The Compute Nodes must have Management Processor (MP) cards.

ClusterPack depends on certain open source software which is normally installed as a part

of the operating environment. The minimum release versions required are:

- MySQL Version 3.23.58 or higher
- Perl Version 5.8 or higher

For more information, see the Comprehensive Instructions for this step.

References:

- [Step 2 Install Prerequisites](#)

[Back to Top](#)

Step Q3 Allocate File System Space

Allocate file system space on the Management Server. Minimum requirements are listed below.

- /var - 4GB
- /opt - 4GB
- /share - 500MB (Clusterware edition only)

For more information, see the Comprehensive Instructions for this step.

References:

- [Step 3 Allocate File System Space](#)

[Back to Top](#)

Step Q4 Obtain a License File

- Get the Host ID number of the Management Server.
- Contact Hewlett-Packard Licensing Services to redeem your license certificates.
- If you purchased the ClusterPack Base Edition, redeem the Base Edition license certificate.
- If you purchased the ClusterPack ClusterWare Edition, redeem the Base Edition certificate and the ClusterWare edition certificate

Note:

It may take up to 24 hours to receive the license file. Plan accordingly.

For more information, see the Comprehensive Instructions for this step.

References:

- [Step 4 Obtain a License File](#)

[Back to Top](#)

Step Q5 Prepare Hardware Access

Get a serial console cable long enough to reach all the Compute Nodes from the Management Server.

Note:

If you are installing ClusterPack on Compute Nodes for the first time, DO NOT power up the systems, ClusterPack will do that for you automatically. If you do accidentally power the compute nodes, DO NOT answer the HP-UX boot questions.

For more information, see the Comprehensive Instructions for this step.

References:

- [Step 5 Prepare Hardware Access](#)

[Back to Top](#)

Step Q6 Power Up the Management Server

Perform a normal boot process for the Management Server.

Note:

If you are using DHCP for IP assignment, DO NOT boot the Compute Nodes at this time.

For more information, see the Comprehensive Instructions for this step.

References:

- [Step 6 Power Up the Management Server](#)

[Back to Top](#)

Step Q7 Configure the ProCurve Switch

- Select an IP address from the same IP subnet that will be used for the Compute Nodes.
- Connect a console to the switch
- Log onto the switch through the console
- Type 'set-up'
- Select IP Config and select the "manual" option
- Select the IP address field and enter the IP address to be used for the switch

For more information, see the Comprehensive Instructions for this step.

References:

- [Step 7 Configure the ProCurve Switch](#)

[Back to Top](#)

Step Q8 Copy the License Files to the Management Server

Put the files in any convenient directory on the Management Server.

For more information, see the Comprehensive Instructions for this step.

References:

- [Step 8 Copy the License Files to the Management Server](#)

[Back to Top](#)

Step Q9 Install ClusterPack on the Management Server

- Mount and register the ClusterPack DVD as a software depot.
- Install the ClusterPack Manager software (CPACK-MGR) using swinstall.
- Leave the DVD in the DVD drive for the next step.

For more information, see the Comprehensive Instructions for this step.

References:

- [Step 9 Install ClusterPack on the Management Server](#)

[Back to Top](#)

Step Q10 Run manager_config on the Management Server

Provide the following information to the manager_config program:

- The path to the license file(s),
- Whether to store passwords,
- The DNS domain and NIS domain for the cluster,
- The host name of the manager and the name of the cluster,
- The cluster LAN interface on the Management Server,
- The count and starting IP address of the Compute Nodes,
- Whether to mount a home directory,
- The SCM admin password,
- The LSF admin password. (Clusterware edition only)

For more information, see the Comprehensive Instructions for this step.

References:

- [Step 10 Run manager_config on the Management Server](#)

[Back to Top](#)

Step Q11 Run mp_register on the Management Server

Provide the following information to the mp_register program about each Management Processor card that is connected to a Compute Node:

- IP address,
- netmask,
- gateway IP address.

For more information, see the Comprehensive Instructions for this step.

References:

- [Step 11 Run mp_register on the Management Server](#)

[Back to Top](#)

Step Q12 Power up the Compute Nodes

Use the clbootnodes program to power up all Compute Nodes that have a connected Management Processor that you specified in the previous step. Provide the following information to the clbootnodes program:

- Language to use,
- Host name,
- Time and time zone settings,
- Network configuration,
- Root password.

For more information, see the Comprehensive Instructions for this step.

References:

- [Step 12 Power up the Compute Nodes](#)

[Back to Top](#)

Step Q13 Run compute_config on the Management Server

The compute_config program will register the nodes with various programs.

For more information, see the Comprehensive Instructions for this step.

References:

- [Step 13 Run compute_config on the Management Server](#)

[Back to Top](#)

Step Q14 Run finalize_config on the Management Server

This program completes the installation and configuration process, verifies the Cluster Management Software, and validates the installation. If it reports diagnostic error messages,

repeat the installation process, performing all steps in the order specified.

For more information, see the Comprehensive Instructions for this step.

References:

- [Step 14 Set up HyperFabric \(optional\)](#)

[Back to Top](#)

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

Copyright 1994-2004 hewlett-packard company

ClusterPack



ClusterPack General Overview

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

[1.1.1 ClusterPack Overview](#)

[1.1.2 Who should use the material in this tutorial?](#)

[1.1.3 What is the best order to review the material in the tutorial?](#)

[1.1.4 Operating System and Operating Environment Requirements](#)

[1.1.5 System Requirements](#)

1.1.1 ClusterPack Overview

Compute Cluster

A cluster is a set of independent computers combined into a unified system through system software and networking technologies. Compute clusters are important options for a growing number of technical and commercial customers. The primary driver for clustering of compute workstations and servers is that it provides a low cost alternative to supercomputers. It also has the following key benefits:

- horizontally scalable by adding more nodes
- vertically scalable by using larger SMP nodes
- fault-isolation - failure of a single Compute Node will not shutdown the entire cluster system
- asymmetry: mix and match of different nodes in a cluster
- configuration flexibility: nodes, interconnect
- re-deployable nodes

A compute cluster consists of Compute Nodes that incorporate multiple processors in a symmetric multiprocessor configuration. These nodes are connected through interconnection network(s) and are supervised within a single administrative domain. A compute cluster usually resides in a single machine room and is managed as a single computer system. The interconnection network employs standard local area network (LAN) and/or high performance interconnect technologies based on the application requirements. A compute cluster has to support both time-to-solution jobs and throughput jobs. HP high performance technical computing (hptc) cluster is available on HP Integrity servers with interconnection

options of Gigabit Ethernet or HyperFabric2. The common components of a cluster are:

- Head Node - provides user access to the cluster. In smaller clusters, the Head Node may also serve as a Management Server.
- Management Server - server that provides single point of management for all system components in the cluster
- cluster LAN/switch - usually an Ethernet network used to monitor and control all the major system components. May also handle traffic to the file server.
- interconnect switch - provides high speed connectivity between Compute Nodes. Used for message passing and remote memory access capabilities for parallel applications.
- Storage - includes both local disk space on each Compute Node and optionally an independent storage unit accessible to the Compute Nodes (See NAS below).
- Cluster management software - ClusterPack for system administrators and end-users.
- Management Processor (MP) - controls the system console, reset and power management functions of a server.
- Network Attached Storage (NAS) - attach directly to Ethernet networks, providing easy installation, low maintenance, and high uptime.

Compute Nodes

Compute Nodes in a cluster are normal compute servers that provide system computing resource and storage capability. An hptc cluster is built with HP Integrity servers: 2-way or 4-way server platforms based on Intel Itanium 2-based processors and HP's zx1 chipset technologies.

The HP Integrity rx2600 server, powered by Intel Itanium 2-based processors, is the industry's first dual-processor Itanium 2-based server. The rx2600 dramatically improves price/performance for technical applications and gives customers a cost-effective yet powerful entry into a technology with tremendous growth potential. Its 2U rack density is ideal for clustering solutions.

HP's extensive Itanium 2-based systems experience and co-developer insights have resulted in premium performance through the development of the HP Scalable Processor Chipset zx1. Invented by HP, the HP zx1 chipset fully unleashes the power of Intel Itanium 2-based processors by lowering memory latencies and increasing memory and I/O subsystem scalability. With the HP zx1 chipset, HP Integrity servers and Itanium 2-based servers achieve even better performance and memory expandability.

Interconnection Networks

An hptc cluster is built with the industry standard Gigabit Ethernet. This interconnect network is capable of transferring data packets between Compute Nodes for both file serving and inter-node communication for applications. When applications benefit from lower latency, HP's HyperFabric2 can be configured into the cluster to provide lower latency and higher bandwidth. Since HyperFabric2 uses the Hyper Messaging Protocol (HMP) that by-passes the TCP/IP stack, CPU usage is also greatly reduced when

transferring data between nodes. A cluster LAN is also configured to separate the system management traffic from application message passing and file serving traffics.

Management Software and Head Node

The ability to manage and use a cluster as easily as a single compute system is critical to the success of any cluster solution. To facilitate ease of use for both system administrators and end-users, HP has created a software package called ClusterPack. The hptc cluster can be managed and used just as easily as a single server platform.

The ClusterPack is an integrated solution that offers the following key features:

Installation and configuration

- automated cluster setup
- network services setup (NFS, NTP, NIS, Ignite-UX)
- remote power-on
- network configuration

System administration

- single point of administration
- role-based management
- software and patch management
- user account maintenance
- resource control
- cluster health monitoring
- cluster troubleshooting
- cluster tuning
- golden image creation and distribution
- cluster reconfiguration
- cluster system hardware and software inventory management
- cluster server nodes consistency checking

Distributed resource management

- cluster resource scheduling
- policy based queues and multiple queue management
- job submission, monitor, and control
- user specified job priority
- Application ReStart file migration

The ClusterPack integrates HP's enterprise system management solution, ServiceControl Manager (SCM), and Platform Computing's latest offering, . HP's SCM has been widely used in large-scale data centers for managing hundreds of HP systems through a single point of control. Its multi-system management capabilities, such as group operations and role-based management, enable customers to achieve optimal IT resource efficiency. Platform Computing's Clusterware is based on the industry distributed resource management solution, Load Sharing Facility (LSF), and is specially designed for cluster management and ease-of-use for both system administrators and end-users. The ClusterPack is available on HP-UX

11i Version 2.0. The ClusterPack has a server component that runs on a Management Server, and client agents that run on the managed Integrity compute servers.

NAS 8000

NAS 8000 High Availability Cluster was designed to significantly reduce downtime and maximize the availability of storage by providing heterogeneous file-sharing and file-serving functionality across a wide variety of application areas, including content delivery and distribution, consolidated storage management, technical computing, and Web serving. It allows capacity without server downtime so additions can be made without losing access to valuable data. The HP NAS 8000 series provides interoperability with other applications. It also provides storage flexibility.

The critical need for more accessible data, and more reliable storage, is met with a NAS solution.

[Back to Top](#)

1.1.2 Who should use the material in this tutorial?

This tutorial is divided into four main sections:

- Administrators Guide
- Users Guide
- Tool Overview
- Related Documents
- Data Dictionary

The Administrators Guide is intended for the system administrators who will be responsible for the initial setup and continuing operation of the cluster. The Administrators section of the tutorial covers a range of topics including: installation and setup of the ClusterPack software on the cluster, creating and managing golden images, system maintenance tasks, adding users to the cluster, adding third party software to the cluster, system monitoring tasks, and basic troubleshooting of the ClusterPack software.

The Users Guide is intended for users who will be using the ClusterPack software to submit jobs to the cluster. The Users section of the tutorial covers a range of topics including: basic user account setup and access, the layout of the cluster, submitting jobs to the cluster, monitoring jobs currently running on the cluster, collecting data in a central location, and general troubleshooting of the ClusterPack software.

The Tool Overview is a general overview of each major software component of the ClusterPack solution. For each tool, a basic functional overview is presented. This section includes information that is useful for both system administrators and end users.

The Related Documents gives the location of additional information for components of ClusterPack and HP Integrity clusters.

The Data Dictionary contains definitions for common terms that are used through the tutorial.

[Back to Top](#)

1.1.3 What is the best order to review the material in the tutorial?

System Administrators

Initial installation and configuration of the cluster requires a complete understanding of the steps involved and the information required. Before installing a new cluster, the system administrator should read and understand all of the steps involved before beginning the actual installation. It is assumed that anyone setting up a cluster is familiar with Unix system administration, HP-UX OS installation, the Management Processor interface and network setup and configuration. Before installing a cluster for the first time, read and understand the following material:

- Section 1.2 Initial Installation and Set-Up
- Section 1.3 Installation and Configuration of Optional Components
- Section 1.4 Software Upgrades and Reinstalls
- Section 1.5 Golden Image Tasks

It is helpful prior to installation to review and be familiar with several additional sections of the tutorial. This material does not need to be completely reviewed, but should be read and available during the initial testing of the new cluster:

- Section 1.6 System Maintenance Tasks
- Section 1.7 System Monitoring Tasks
- Section 1.8 Workload Management Tasks
- Section 1.9 System Troubleshooting Tasks
- Tool Overview

Users

Before beginning to use the cluster for daily tasks, it is helpful to review several sections in order to be familiar with the tools that are provided on the cluster to submit and monitor jobs. This material should be kept readily available for reference.

- Section 2.1 Job Management Tasks
- Section 2.2 File Transfer Tasks
- Section 2.3 Miscellaneous Tasks
- Tool Overview

Is the Online Tutorial available in a single printable file?

The Online Tutorial is available in a single file. From the Master Index of Sections, there is

a link to the printable version at the bottom of the page.

References:

- [Printable Version](#)

[Back to Top](#)

1.1.4 Operating System and Operating Environment Requirements

The key components of the HP Integrity Server Technical Cluster are:

- Management Server: HP Integrity server with HP-UX 11i Version 2.0 TCOE
- Compute Nodes: HP Integrity servers with HP-UX 11i Version 2.0 TCOE
- Cluster Management Software: ClusterPack V2.3

The following prerequisites are assumed:

- HP-UX 11i V2.0 TCOE installed on the Management Server
- HP-UX 11i V2.0 TCOE installed on each Compute Node

The following software components must be installed for all features of ClusterPack V2.3 to function effectively:

- HP-UX 11i Ignite-UX on the Management Server
- HP-UX 11i Ignite-UX on each Compute Node

[Back to Top](#)

1.1.5 System Requirements

In order to install the ClusterPack solution, the Management Server must have access to a DVD drive.

The Management Server hard drive must be partitioned to allow sufficient disc space for the following directories:

- /var - 4 GB
- /opt - 4 GB
- /share - 500 MB

[Back to Top](#)

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

Copyright 1994-2004 hewlett-packard company

Comprehensive Install Instructions

ClusterPack



Comprehensive Install Instructions

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

[1.2.1 Comprehensive Installation Overview](#)

[Step 1 Fill Out the ClusterPack Installation Worksheet](#)

[Step 2 Install Prerequisites](#)

[Step 3 Allocate File System Space](#)

[Step 4 Obtain a License File](#)

[Step 5 Prepare Hardware Access](#)

[Step 6 Power Up the Management Server](#)

[Step 7 Configure the ProCurve Switch](#)

[Step 8 Copy the License Files to the Management Server](#)

[Step 9 Install ClusterPack on the Management Server](#)

[Step 10 Run manager_config on the Management Server](#)

[Step 11 Run mp_register on the Management Server](#)

[Step 12 Power up the Compute Nodes](#)

[Step 13 Run compute_config on the Management Server](#)

[Step 14 Set up HyperFabric \(optional\)](#)

[Step 15 Set up InfiniBand \(optional\)](#)

[Step 16 Run finalize_config on the Management Server](#)

[Step 17 Create a Golden Image of a Compute Node from the Management Server](#)

[Step 18 Add nodes to the cluster that will receive the Golden Image](#)

[Step 19 Distribute the Golden Image to the remaining Compute Nodes](#)

[Step 20 Install and Configure the remaining Compute Nodes](#)

[Step 21 Verify the final cluster configuration](#)

1.2.1 Comprehensive Installation Overview

ClusterPack uses a two-stage process for setting up an HP-UX Cluster:

- Create a base configuration with a Management Server and one Compute Node
 - Prepare for installation.
 - Install and configure the Management Server.
 - Install and configure the initial Compute Node and its Management

Processor.

- Verify the Management Server and the initial Compute Node.
- Configure the remaining Compute Nodes with a Golden Image.
 - Create a Golden Image.
 - Add nodes to the configuration that will receive the Golden Image.
 - Distribute the Golden Image to remaining nodes.
 - Install and configure the Compute Nodes that received the Golden Image.
 - Verify the final cluster configuration.

These processes are further broken down into a number of discrete steps. Each step contains the following sections:

- Background
- Overview
- Details

The background section explains why this step is necessary and what will be done for you. The overview section tells you what this step entails in general terms. The details section gives the exact commands you must enter.

Note:

The steps in this section have to be followed in the specified order to ensure that everything works correctly. Please read all of the following steps BEFORE beginning the installation process.

[Back to Top](#)

Step 1 Fill Out the ClusterPack Installation Worksheet

Background

ClusterPack simplifies the creation and administration of a cluster of HP Integrity Servers running HP-UX by automating the collection, recording, and distribution of information about the systems in a network. The system administrator, must still make decisions about how to identify and secure those network components. You can record all of your decisions on a form that you can keep nearby as you execute the steps of the installation process.

Overview

Print out this form and fill out all information for each node in your cluster.

[Installation Worksheet \(pdf\)](#)

Note:

You will not be able to complete the following steps if you have not collected all of this information.

Details

At various points during the configuration you will be queried for the following information:

- DNS Domain name [ex. domain.com]
- NIS Domain name [ex. hpcluster]
- Network Connectivity:
 - Information on which network cards in each Compute Node connect to the Management Server
 - Information on which network card in the Management Server connects to the Compute Nodes.
- SCM Administrator password (You will be asked to set it).
- LSF Administrator password (You will be asked to set it).

[Back to Top](#)

Step 2 Install Prerequisites

Background

ClusterPack works on HP Integrity Servers running HP-UX. In order to install ClusterPack, you must have the Technical Computing Operating Environment (TCOE) version of HP-UX installed. You must also have the Ignite-UX software, which is used for installation. Installing Ignite-UX on the Compute Nodes makes it possible to create and distribute "golden images" from the Compute Nodes.

ClusterPack requires a homogeneous operating system environment. That is, all Compute Nodes and the Management Server must have the same release of HP-UX installed as well as the same operating environment.

Overview

Install the following software on the Management Server.

- HP-UX 11i Ignite-UX
- HP-UX 11i V2.0 TCOE

Install the following software on each Compute Node.

- HP-UX 11i Ignite-UX

- HP-UX 11i V2.0 TCOE

ClusterPack depends on certain open source software which is normally installed as a part of the operating environment. The minimum release versions required are:

- MySQL Version 3.23.58 or higher
- Perl Version 5.8 or higher

The Management Server requires a minimum of two LAN connections. One connection must be configured prior to installing ClusterPack.

The Compute Nodes must have Management Processor (MP) cards.

Details

Install these items when you do a fresh install of HP-UX. on the Management Server and the Compute Nodes. Or, you can install Ignite-UX after rebooting by the following method.

- Using the HP-UX 11i V2.0 TCOE DVD, mount and register the DVD as a software depot.
- Install the Ignite-UX software on the Management Server using swinstall.

On the Management Server:

```
% /usr/sbin/swinstall -s <source_machine>:/mnt/dvdrom Ignite-UX
```

Note:

Allow the default choices to install.

[Back to Top](#)

Step 3 Allocate File System Space

Background

ClusterPack installs software in the /opt and /share file systems. It stores data in the /var file system. You must allocate sufficient space in these file systems if it is to operate correctly.

Overview

Allocate file system space on the Management Server. Minimum requirements are listed below.

- /opt - 4GB
- /var - 4GB

- /share - 500MB (Clusterware edition only)

Details

Allocate space for these file systems when you do a fresh install of HP-UX on the Management Server.

To resize /opt

1. Go to single user mode.

```
% # /usr/sbin/shutdown -r now
```

2. Interrupt auto boot.
3. Select the EFI shell.
4. Select the appropriate file system. (Should be fs0: but may be fs1:)

```
% Shell> fs0:
```

5. Boot HP-UX.

```
% fs0:\>hpx
```

6. Interrupt auto boot.
7. Boot to single user mode.

```
% HPUX> boot vmunix -is
```

8. Determine the lvol of /opt.

```
% cat /etc/fstab
```

9. Look for the lvol that corresponds to /opt.
10. Extend the file system. (Use lvol from Step 2.)

```
% # lvextend -L 4096 /dev/vg00/lvol4
```

(May not be lvol4.)

```
% # umount /dev/vg00/lvol4
```

(This should fail.)

```
% # extendfs /dev/vg00/lvol4
```

```
% # mount /dev/vg00/lvol4
```

11. Repeat 8 through 10 for /var.

[Back to Top](#)

Step 4 Obtain a License File

Background

For ClusterPack Base Edition, please refer to the Base Edition License certificate for instructions on redeeming your license.

For ClusterPack Clusterware Edition, you will need to redeem BOTH the Base Edition license certificate AND the Clusterware Edition license certificate. You will need TWO license files in order to run manager_config.

As part of the normal installation and configuration process, you will be asked to provide the license key. ClusterPack will install the license files in the correct location(s), and all licensing services will be started.

Overview

- Get the Host ID number of the Management Server.
- Contact Hewlett-Packard Licensing Services to redeem your license certificates.
- If you purchased the ClusterPack Base Edition, redeem the Base Edition license certificate.
- If you purchased the ClusterPack Clusterware Edition, redeem the Base Edition certificate and the Clusterware edition certificate

Note:

It may take up to 24 hours to receive license file. Plan accordingly.

Details

You will need to contact HP licensing to redeem your license certificates. You can call, E-mail, or fax your request to Hewlett-Packard Software Licensing Services. Refer to your Software License Certificate for contact information. Prior to installing ClusterPack V2.3, you can request a key by providing the Host ID number of the Management Server. The Host ID can be found using the 'uname' command:

```
% /bin/uname -i
```

The number returned by this command must be proceeded by a # when making your request. For example, 'uname -i' returns 2005771344, provide the number as #2005771344 in your key request.

[Back to Top](#)

Step 5 Prepare Hardware Access

Background

This document does not cover hardware details. It is necessary, however, to make certain hardware preparations in order to run the software.

Overview

Get a serial console cable long enough to reach all the Compute Nodes from the Management Server.

Details

To allow the Management Server to aid in configuring the Management Processors, it is necessary to have a serial console cable to connect the serial port on the Management Server to the console port on the Management Processor to be configured. Be sure that the serial cable is long enough to reach all of the Compute Nodes. It is also possible to configure the Management Processors manually by connecting a console to each card.

Note:

If you are installing ClusterPack on Compute Nodes for the first time, DO NOT power up the systems, ClusterPack will do that for you automatically. If you do accidentally power the compute nodes, DO NOT answer the HP-UX boot questions.

[Back to Top](#)

Step 6 Power Up the Management Server

Background

This is the first step in actually configuring your system.

Overview

Perform a normal boot process for the Management Server.

Note:

If you are using DHCP for IP assignment, DO NOT boot the Compute Nodes at this time.

Details

If you use DHCP to assign the host names and IP addresses to the Compute Nodes, the DHCP server must be running on the Management Server before the Compute Nodes are booted. The manager_config program will set up the DHCP server.

```
% /opt/clusterpack/bin/manager_config
```

[Back to Top](#)

Step 7 Configure the ProCurve Switch

Background

The ProCurve Switch is used for the management network of the cluster.

Overview

The IP address for the ProCurve Switch should be selected from the same IP subnet that will be used for Compute Nodes.

Details

- Select an IP address from the same IP subnet that will be used for the Compute Nodes.
- Connect a console to the switch
- Log onto the switch through the console
- Type 'set-up'
- Select IP Config and select the "manual" option
- Select the IP address field and enter the IP address to be used for the switch

[Back to Top](#)

Step 8 Copy the License Files to the Management Server

Background

Copy the license files onto the Management Server. The license files can be placed in any convenient directory that is accessible to the Management Server. During the invocation of the manager_config tool you will be asked to provide a path to the license files. As part of manager_config the license files will be installed into the correct locations on the machine, and all licensing services will be started.

Overview

Put the files in any convenient directory on the Management Server.

Details

```
% /usr/bin/ftp your_host
% > cd your_home
% > lcd /tmp
% > get cpack.lic
% > bye
```

[Back to Top](#)

Step 9 Install ClusterPack on the Management Server

Background

The ClusterPack software is delivered on a DVD.

Overview

- Mount and register the ClusterPack DVD as a software depot.
- Install the ClusterPack Manager software (CPACK-MGR) using swinstall.
- Leave the DVD in the DVD drive for the next step.

Details

How to mount a DVD on a remote system to a local directory

On the system with the DVD drive (i.e. remote system):

1. Mount the DVD.

```
% mount /dev/dsk/xxxx /mnt/dvdrom
```

2. Edit the /etc/exports file. DVDs must be mounted read only ('ro'), and if required, can give root permission to other machines mounting the filesystem ("root=<machine_foo:machine_bar:machine_baz>"). Add a line to /etc/exports:

```
% /mnt/dvdrom -ro,root=<local_system>
```

3. Export the file system, using all the directives found in /etc/exports

```
% exportfs -a
```

4. Check to be sure that the line you added is actually exported.

```
% exportfs
```

On the local machine:

5. Mount the DVD to an existing directory.

```
% /etc/mount <remote_system>:/mnt/dvdrom /mnt/dvdrom
```

Note:

You cannot be in the /mnt/dvdrom directory when you try to mount. You will get a file busy error.

When you are finished, on the local machine:

6. Unmount the DVD file system.

```
% /etc/umount /mnt/dvdrom
```

On the remote system:

7. Unexport the DVD file system.

```
% exportfs -u -i /mnt/dvdrom
```

8. Unmount the DVD

```
% /etc/umount /mnt/dvdrom
```

How to enable a DVD as a software depot

During the installation process, two DVDs will be required. Generic instructions for making a DVD accessible as a software depot for installation onto the Management Server are provided here. Please refer to the steps that follow for the specific DVDs that are required.

The steps to mount a DVD for use as a software depot are:

- Insert DVD into the drive
- Mount the DVD drive locally on that system
- Register the depot on the DVD using swreg
- Check the contents of the DVD using swlist

These commands can only be executed as the super user (i.e. root).

A DVD drive installed in the Management Server can be used for software installations. If the Management Server does not include a DVD drive, use one of these two methods:

1. Connect a portable DVD drive to the Management Server
2. Use a HP-UX system with a DVD drive, that is network accessible from the Management Server, as a source for installation.

For example, to mount the device /dev/dvdrom to the directory /mnt/dvdrom, execute the following commands on the "source machine" with the DVD drive

```
% /sbin/mount -r /dev/dsk/xxxx /mnt/dvdrom
% /usr/sbin/swreg -l depot /mnt/dvdrom
% /usr/sbin/swlist @ /mnt/dvdrom
```

- Using the ClusterPack DVD, mount and register the DVD as a software depot.
- Install the ClusterPack Manager software (CPACK-MGR) on the Management Server using swinstall.

On the Management Server:

```
% /usr/sbin/swinstall -s <source_machine>:/mnt/dvdrom CPACK-MGR
```

- The ClusterPack DVD will be referenced again in the installation process. Please leave it in the DVD drive until the "Invoke /opt/clusterpack/bin/manager_config on Management Server" step has completed.

[Back to Top](#)

Step 10 Run manager_config on the Management Server

Background

This program is the main installation and configuration driver. It should be executed on the Management Server.

Some of the steps are:

- Install the appropriate license files (i.e. ONE Base Edition license file or TWO Clusterware Edition license files), and start the licensing services
- Assign DNS domain name and NIS domain name based on inputs provided
- Select and configure the cluster LAN interface on the Management Server that interfaces with the Compute Nodes
- Specify how many Compute Nodes are in the cluster and the starting IP address of the first Compute Node. This information is used to assign names and IP addresses when Compute Nodes are brought up. The first 5 characters of the Management Server's hostname are used for a base for the Compute Nodes. For example, if the starting IP address is 10.1.1.1, and there are 16 Compute Nodes, and the name of the Management Server is hpnode, then the first Compute Node will be called hpnod001 with address 10.1.1.1, the next hpnod002 with address 10.1.1.2, and so on. If the tool is invoked with the -f option, the input file will be the source for this information.
- Set up the Management Server as NTP server, NIS server, NFS server, Ignite-UX server, and Web server.
- Install all of the dependent software components from the ClusterPack DVD:
 - This step looks for the source of the CPACK-MGR install and queries for an alternate source, if the source is not found. A local depot is setup. All of the agent components are copied. Other dependent software pieces in the Management Server are validated and installed.
- Configure the Management Server components of SCM, SIM, and other Cluster Management Software tools.
- Modify configuration files on the Management Server to enable auto-startup of the Cluster Management Software components after reboots.
- Print a PASS diagnostic message if all of the configuration steps are successful.

Overview

Provide the following information to the manager_config program:

- The path to the license file(s),
- Whether to store passwords,
- The DNS domain and NIS domain for the cluster,
- The host name of the manager and the name of the cluster,
- The cluster LAN interface on the Management Server,
- The count and starting IP address of the Compute Nodes,
- Whether to mount a home directory,
- The SCM admin password,
- The LSF admin password. (Clusterware edition only)

Details

This tool can be invoked in two ways, based on your specific requirements:

- If you want manager_config to drive the allocation of hostnames and IP addresses of the Compute Nodes in the cluster (based on some basic queries), /opt/clusterpack/bin/manager_config is invoked with no arguments:

```
% /opt/clusterpack/bin/manager_config
```

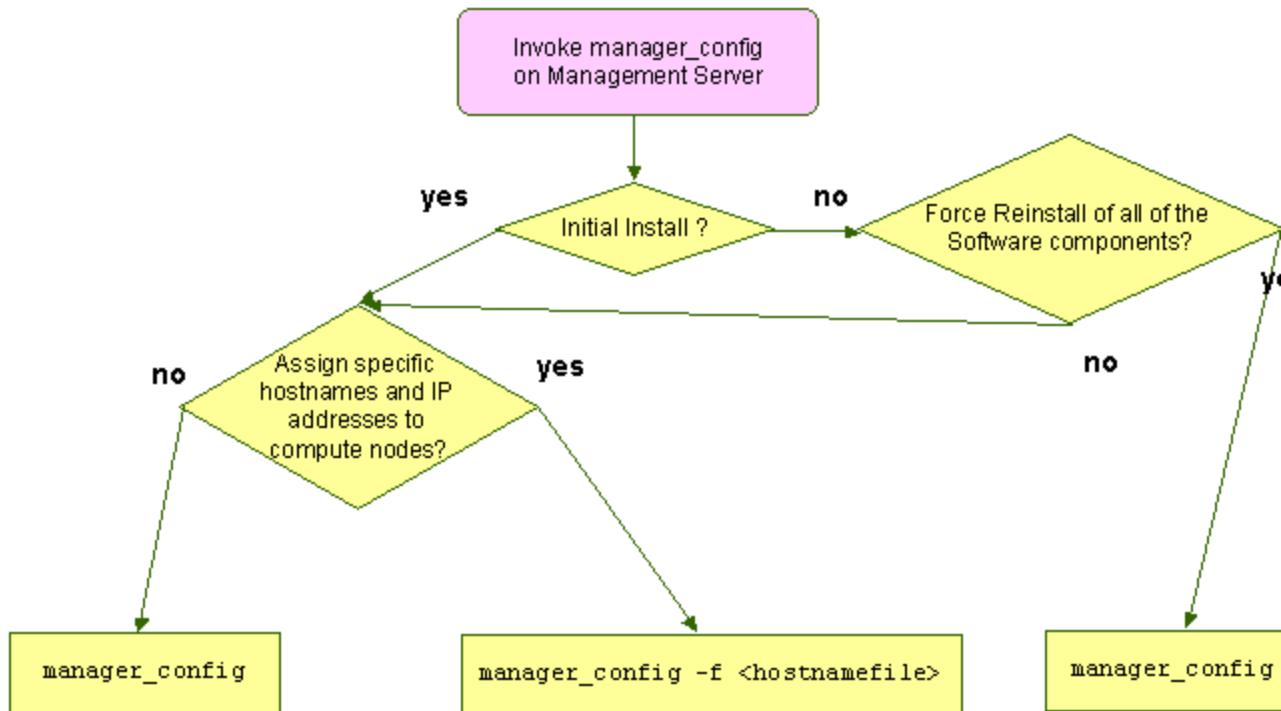
- If you want manager_config to assign specific hostnames and IP addresses to the Compute Nodes in the cluster, supply an input file in the same format as /etc/hosts, and invoke the tool as follows:

```
% /opt/clusterpack/bin/manager_config -f input_file
```

- The ClusterPack DVD is no longer required during installation. On the source machine, unmount the DVD drive and remove the DVD.

```
% /sbin/umount /mnt/dvdrom
```

manager_config Invocation



manager_config Invocation

manager_config is an interactive tool that configures the Management Server based on some simple queries (most of the queries have default values assigned, and you just need to press RETURN to assign those default values).

[Back to Top](#)

Step 11 Run mp_register on the Management Server

Background

A Management Processor (MP) allows you to remotely monitor and control the state of a Compute Node by configuring and registering the MP cards for each Compute Node. clbootnodes can be used to automatically answer the first boot questions for each Compute Node.

When you telnet to an MP, you will initially access the console of the associated server. Other options such as remote console access, power management, remote re-boot operations, and temperature monitoring are available by typing control-B from the console mode. It is also possible to access the MP as a web console. However, before it is possible to access the MP remotely it is first necessary to assign an IP address to each MP. This is normally achieved by connecting a serial console device to the serial port on the MP and performing a series of configuration steps. This can be quite tedious and time consuming for moderate to large clusters. To ease the effort, mp_register can perform the configuration for you by issuing the commands via a serial cable.

mp_register maintains a database of knowledge about the MP cards in the system. The database is restricted to nodes that have been added to the cluster with manager_config. Likewise, nodes removed from the cluster are removed from the MP database. The utility is generally designed for single use when setting up the cluster for the first time. However, it can be run multiple times to make changes to MP designations or when nodes are added to the cluster.

Note:

It is important to note that the configuration step does not configure accounts for the MP. By default, anyone can access the MP without a password. Leaving the cards without configured users is a severe security risk. Users can freely access the card and shut down the node or gain root access through the console. The configuration step configures the MP for telnet or web access only to make future modifications, such as adding users simpler to perform.

mp_register will add each MP and associated IP address to the /etc/hosts file on the Management Server. This file will later get propagated to the Compute Nodes. Each MP is assigned a name during the configuration step which is also placed in the /etc/hosts file. This name is derived as the name of the associated host appended with '-mp' (for Management Processor). For example, the MP associated with the host foo will be named foo-mp.

Overview

Provide the following information to the mp_register program about each MP card that is connected to a Compute Node. It will configure all MPs automatically, instead of requiring you to manually connect the card to a serial console device.

- IP address,
- Netmask,
- Gateway IP address.

Details

For each node, the program will ask you if you want to establish an MP for that machine. It will also ask if the MP is already configured. If it is not already configured, you will be prompted to connect a serial cable from the serial port of the Management Server to the serial port of the MP to be configured. The program will then use the information you entered about the card to configure it. Each MP can be configured in turn. MPs which have been previously configured can be added to the database without being configured.

Before invoking mp_register to initially configure the MP cards on each Compute Node, obtain a serial cable long enough to connect from the serial console port on the back of the Management Server to the serial

console port on the MP card of each Compute Node.

When you are ready to run mp_register, use this command:

```
% /opt/clusterpack/bin/mp_register
```

[Back to Top](#)

Step 12 Power up the Compute Nodes

Background

The clbootnodes utility is intended to ease the task of booting Compute Nodes for the first time. To use clbootnodes, the nodes' MP cards must have been registered and/or configured with mp_register.

Note:

clbootnodes can only be used to boot nodes to the language specification: English

The first time that HP-UX is booted after installation, it asks a series of questions:

- What language to use
- Hostname
- Time and Timezone settings
- Networking Configuration
- Root password

Booting each node in a medium to large cluster can be a long and tedious task. clbootnodes automates the processes to make it much faster and relatively free of user interaction. It is also possible to boot only specified nodes using clbootnodes.

clbootnodes will gain console access by using telnet to reach the MP. clbootnodes uses a library called Expect to produce the input needed to gain access to the console and step through the boot processes. There are times when human intervention is necessary. In these cases, a message will be displayed explaining what control is being returned to the user. The user can then interact with the MP/console and then return control to clbootnodes by pressing '~'. Control may be given to the user for the following reasons:

- The MP is password protected
- A LAN card choice was not specified to clbootnodes
- The utility could not determine the state of the console

clbootnodes is intended to boot a node or nodes through the first boot sequence. It can generally be run at any time to ensure that a node is booted and can usually recognize if the console represents an already booted node. However, because a user can leave the console in any state, it is not always possible to determine the state that a console is in. Because of this, use clbootnodes for booting nodes which are known to be in a "first boot" condition.

When booting a node, clbootnodes will answer the first boot questions rather than having to answer them manually. The questions are answered using the following information:

- Language selection: All language selection options are set to English.
- Keyboard selection: The keyboard selection is US English
- Timezone: The time zone information is determined based on the setting of the Management Server
- Time: The current time is accepted. The time will later be synchronized to the Management Server using NTP.
- Networking: The LAN card specified will be configured to the IP address specified to manager_config.
- Hostname: The hostname will be set to the name specified through manager_config.
- Root password: The root password will be queried before the nodes are booted.

Overview

Use the clbootnodes program to power up all Compute Nodes that have a connected MP that you specified in the previous step. It will answer the first boot questions for all nodes automatically.

Provide the following information to the clbootnodes program:

- Language to use,
- Host name,
- Time and time zone settings,
- Network configuration,
- Root password.

Details

To run clbootnodes, use the following command:

```
% /opt/clusterpack/bin/clbootnodes
```

Before booting the nodes, clbootnodes will ask you for the root password to set on the Compute Nodes and the LAN card to configure for networking for each host. The LAN card choice for each host will be set to the IP address specified earlier via manager_config.

To boot the node using DHCP, invoke the program as follows:

```
% /opt/clusterpack/bin/clbootnodes <list_of_node_names>
```

You can omit the argument list, in which all nodes in the cluster will be processed. The IP address used will be the one that you provided previously. The program will interact with you to obtain the name of the LAN card to use.

[Back to Top](#)

Step 13 Run compute_config on the Management Server

Background

This tool is the driver that installs and configures appropriate components on every Compute Node.

- Registers Compute Nodes with SCM and SIM on the Management Server.
- Pushes agent components to all Compute Nodes.
- Sets up each Compute Node as NTP client, NIS client, and NFS client.
- Starts necessary agents in each of the Compute Nodes.
- Modifies configuration files on all Compute Nodes to enable auto-startup of agents after reboots.
- Allows for the configuration of additional networks with `clnetworks`
- Prints a PASS diagnostic message if all configuration steps are successful.

`clnetworks`

Each Compute Node is known to the Management Server through the IP address and specified to `manager_config`. These interfaces are collectively known as the Cluster Network. This term can be somewhat confusing when a cluster consists of both private nodes and public nodes. This is possible, for example, when an initial set of Compute Nodes is created on a private network and then additional nodes outside the private network are added using `-a`. The IP address of each Compute Node known by the Management Server makes up the Cluster Network.

ClusterPack includes a utility to configure additional networks on all of the Compute Nodes. These networks, like the Cluster Network, refer to a logical collection of interfaces/IP addresses and not to a physical network. However, they must share a common netmask. The concept of a network is defined as:

- A name (for reference only)
- A subset of the nodes in the cluster
- A network interface for each node in the subset
- An IP address for each interface
- A name extension that is added to the hostname of each machine and associated with each host's interface.
- A netmask

To define additional networks, use the command `clnetworks`. This tool is also called from `compute_config`.

`clnetworks` provides a text-based interface for selecting nodes, network interfaces and IP addresses. It guides the user through the creation of a network. It is also possible to modify an existing network. When you have finished creating or updating networks, `clnetworks` will ensure that each interface specified is configured correctly and ensure the proper entries in each host's `/etc/hosts` file.

Overview

The `compute_config` program will register the nodes with various programs.

Details

Execute the following command.

```
% /opt/clusterpack/bin/compute_config
```

[Back to Top](#)

Step 14 Set up HyperFabric (optional)

Background

The utility clnetworks assists in setting up a HyperFabric network within a cluster. For clnetworks to recognize the HyperFabric (clic) interface, it is necessary to first install the drivers and/or kernel patches are needed.

Once the clic interface is recognized by lanscan, clnetworks can be used to set (or change) the IP address configure the card.

The HyperFabric interfaces can be grouped into a network entity using clnetworks. The network entity m be assigned an extension that forms the aliases to use for the HyperFabric interfaces. Use these names w\ you want to explicitly communicate over the HyperFabric network. For example, if node002 has a HyperFabric interface with the extension "hyp", ftp through this network can be achieved using:

```
% /usr/bin/ftp node002-hyp
```

Notice that this command will only work from a node that has a HyperFabric interface.

Overview

Install the HyperFabric drivers and kernel patches.

Details

Use the swinstall command to install the following bundle:

- HyprFabrc-00

Follow the installation instructions delivered to you with your HyperFabric installation kit.

[Back to Top](#)

Step 15 Set up InfiniBand (optional)

Background

ClusterPack configures IP over InfiniBand (IPoIB) if the appropriate InfiniBand drivers are installed on t

systems.

Overview

If the InfiniBand IPoIB drivers are installed prior to running `compute_config`, the InfiniBand HCA is detected and the administrator is given a chance to configure them.

The administrator can also configure the InfiniBand HCA with IP addresses by invoking `/opt/clusterpack/bin/clnetworks`. See the man pages for `clnetworks` for usage instructions.

[Back to Top](#)

Step 16 Run `finalize_config` on the Management Server

Background

This step performs verification checks on the Cluster Management Software, and validates the installation. It prints out diagnostic error messages if the installation is not successful.

Note:

The `finalize_config` tool can be run at any time to validate the cluster configuration and to determine if there are any errors in the ClusterPack software suite.

Overview

This program verifies the Cluster Management Software, and validates the installation of the single Compute Node. If it reports diagnostic error messages, repeat the installation process up to this point, performing a few steps in the order specified.

Details

Finalize and validate the installation and configuration of the ClusterPack software.

```
% /opt/clusterpack/bin/finalize_config
```

[Back to Top](#)

Step 17 Create a Golden Image of a Compute Node from the Management Server

Background

A system image is an archive of a computer's file system. Capturing the file system of a computer captures the state of the computer at the time of capture.

the basic state of a computer system. An image does not generally include all files however. By default, / and other temporary files, network directories and host specific configuration files are not included.

A system image may be referred to as a golden image or a recovery image. The different names used to refer to the image reflect the different reasons for creating it. Administrators may create a "recovery" image of a node in the event that the node experiences hardware failure or the file system is accidentally removed or corrupted. Administrators may also create a "golden" image for the purpose of installing it on other nodes to insure that each node in their cluster is configured exactly the way they want.

Overview

- Clean up anything on the system that shouldn't be in the image.
- Ensure that the system isn't being used.
- Run sysimage_create to create the Golden Image.

Details

- Log into the Compute Node to be archived.
- Perform general file system cleanup/maintenance. For example, it may be desirable to search for and remove core files.

Next, from the Management Server:

- Ensure that the system is not being used. It is advisable that the system stop accepting new LSF jobs while the archive is being made:

```
% badmin hclose <hostname>
```

- In addition, you should either wait until all running jobs complete, or suspend them:

```
% bstop -a -u all -m <hostname>
```

- Execute sysimage_create on the Management Server and pass the name of the file from which you would like the image to be made. For example:

```
% /opt/clusterpack/bin/sysimage_create <hostname>
```

- Monitor the output for possible error conditions. The image will be stored in /var/opt/ignite/archives/<hostname>

After the image creation is complete, the Compute Node should be opened for accepting Clusterware job

```
% badmin hopen <hostname>
```

[Back to Top](#)

Step 18 Add nodes to the cluster that will receive the Golden Image

Background

This command adds the new node with the specified host name and IP address to the cluster. It also reconfigures all of the components of ClusterPack to accommodate the newly added node.

Details

Invoke /opt/clusterpack/bin/manager_config with the "add node" option (-a). You can include multiple host:ip pairs if you need to.

```
% /opt/clusterpack/bin/manager_config -a <new_node_name>:<new_node_ip_addr>
```

[Back to Top](#)

Step 19 Distribute the Golden Image to the remaining Compute Nodes

Background

This is the step that actually installs the Golden Image on the Compute Nodes.

Overview

- Register the image.
- Distribute the image to selected nodes.

Details

To distribute a Golden Image to a set of Compute Nodes, you need to first register the image. To register image, use the command:

```
% /opt/clusterpack/bin/sysimage_register <full path of image>
```

If the image was created with sysimage_create, the full path of the image was displayed by sysimage_create. Images are stored in the directory:

```
/var/opt/ignite/archives/<hostname>
```

To distribute the Golden Image to the Compute Nodes, use the command:

```
% /opt/clusterpack/bin/sysimage_distribute <full path of image> [ hostname | all ]
```

The keyword "all" can be used to distribute the image to all of the Compute Nodes in the cluster, or a single hostname can be specified. sysimage_distribute will reboot each Compute Node for installation with the specified image.

[Back to Top](#)

Step 20 Install and Configure the remaining Compute Nodes

Background

This tool is the driver that installs and configures appropriate components on every Compute Node.

Overview

Perform this process in the same way as configuring the first Compute Node.

References:

-

Details

Use the following command to install and configure a Compute Node that received the Golden Image. Perform this for all nodes. You can specify multiple nodes on the command line. You must place the option "-a" in front of each node name.

```
% /opt/clusterpack/bin/compute_config -a <node name>
```

[Back to Top](#)

Step 21 Verify the final cluster configuration

Background

This step completes the installation and configuration process, performs verification checks on the Cluster Management Software, and validates the installation. It prints out diagnostic error messages if the installation is not successful.

Overview

This program completes the installation and configuration process, verifies the Cluster Management Software, and validates the installation. If it reports diagnostic error messages, repeat the installation process performing all steps in the order specified.

Details

Finalize and validate the installation and configuration of the ClusterPack software.

```
% /opt/clusterpack/bin/finalize_config
```

[Back to Top](#)

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

Copyright 1994-2004 hewlett-packard company

ClusterPack



Installation and Configuration of Optional Components

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

[1.3.1 HP-UX IPFilter](#)

[1.3.2 External /home File Server](#)

[1.3.3 Adding Head Nodes to an ClusterPack cluster](#)

[1.3.4 Set up TCP-CONTROL](#)

1.3.1 HP-UX IPFilter

Introduction to NAT (Network Address Translation)

Network Address Translation (NAT) or IP aliasing provides a mechanism to configure multiple IP addresses in the cluster to be presented as a single image view with a single external IP address.

The importance of this is two-fold:

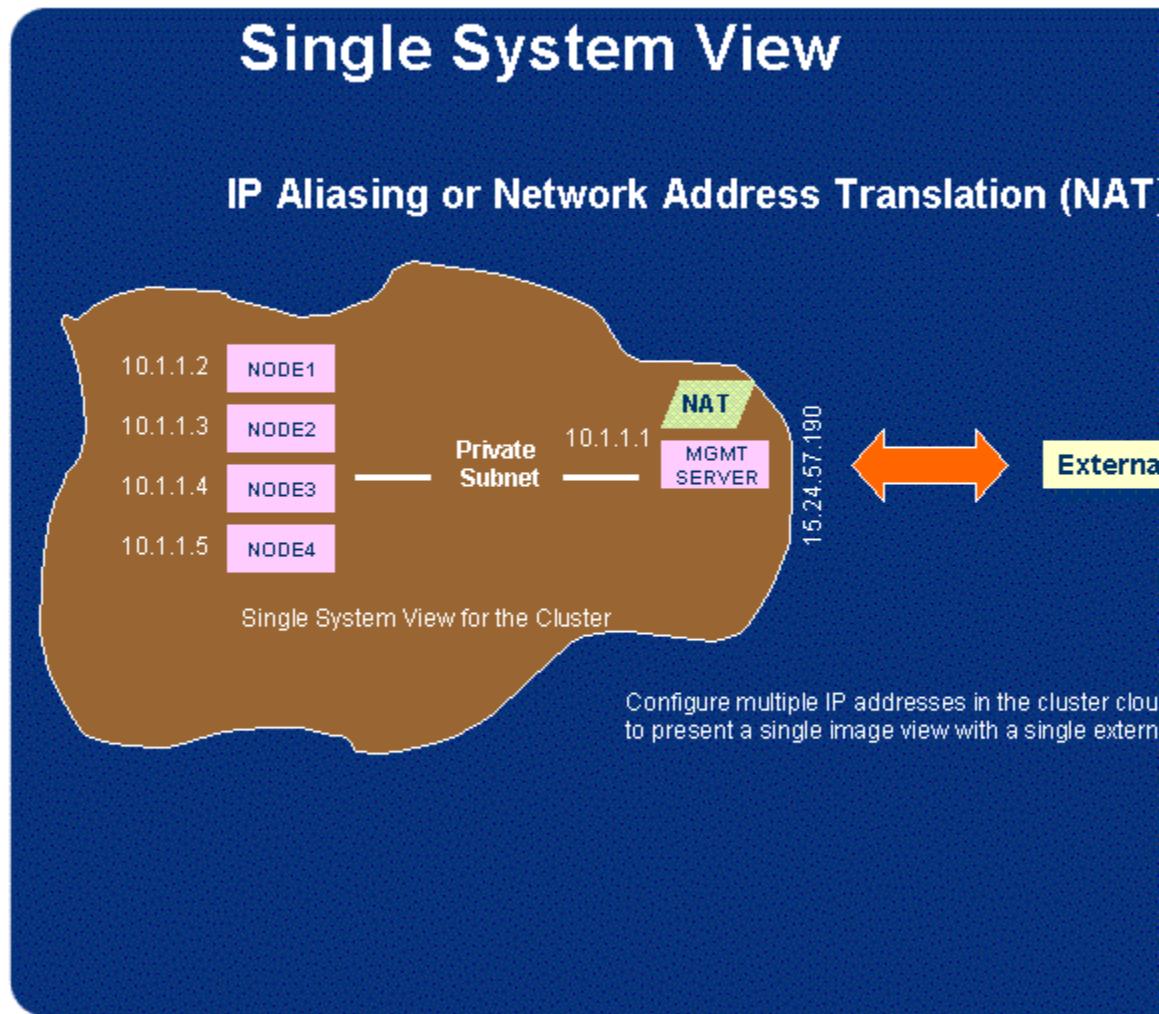
1. Single point of "controlled" access to all of the cluster nodes.

Network Address Translation allows communications from inside the cluster to get out, without allowing connections from outside to get in. NAT rewrites the IP headers of internal packets going out, making it appear that they all came from a single IP address (which is the external IP address of the entire cluster). Reply packets coming back are translated back, and forwarded to the appropriate Compute Node. Thus, the Compute Nodes are allowed to connect to the outside world, if needed. However, outside machines cannot initiate any connection to individual Compute Nodes, since they are exposed only to the "translated" IP address of the entire cluster.

2. Simplified network administration.

The administrator can add or delete nodes to/from the cluster Compute Nodes without any impact to the external world. Further, the administrator can have the all the Compute

Nodes in a private IP sub-net (10.x.y.z range, 192.168.p.q range), which also alleviates the need for numerous public IP addresses.



IP Aliasing or Network Address Translation (NAT)

ClusterPack comes with HP-UX IPFilter, a software component with powerful packet filtering and firewalling capabilities. One of the features that it supports is Network Address Translation.

For your information on HP-UX IPFilter, please refer to the HP-UX IPFilter manual and release notes at docs.hp.com:

<http://docs.hp.com/hpux/internet/index.html#IPFilter/9000>

For information on NAT features of HP-UX IPFilter refer to the public domain how-to document. No guarantee can be made about the correctness, completeness or applicability of this or any third party information.

<http://www.obfuscation.org/ipf/>

HP-UX IPFilter Validation

HP-UX IPFilter is installed with the default HP-UX 11i V2 TCOE bundle. To validate its installation, run the following command:

```
% swverify B9901AA
```

Automatic setup of HP-UX IPFilter rules

ClusterPack V2.3 provides a utility called nat.server to automatically set up the NAT rules, based on the cluster configuration. This tool can be invoked as follows:

```
% /opt/clusterpack/lbin/nat.server
```

Manually setting up the HP-UX IPFilter and NAT rules

Setup the HP-UX IPFilter filtering rules

In this section, we will walk through the steps of setting up HP-UX IPFilter pass through all of the packets. For more complicated filtering rules, please refer to the HP-UX IPFilter documentation.

- Create a file with pass-through rules

```
% cat << EOF > /tmp/filter.rules
pass in all
pass out all
EOF

% cat /tmp/filter.rules
pass in all
pass out all
```

To create more complicated rules, please refer to the HP-UX IPFilter documentation.

<http://docs.hp.com/hpux/internet/index.html#IPFilter/9000>

- Enable filtering based on this rule set

```
% ipf -f /tmp/filter.rules
```

Note:

If there is an existing rule set, a new rule set can be loaded using:

```
% ipf -Fa -f /tmp/filter.rules
```

- For more complicated manipulations of the rule sets, refer to the ipf(5) and ipf(8) man pages.

```
% man 5 ipf
```

```
% man 8 ipf
```

- List the input output filter rules

```
% ipfstat -hio
```

Setup the NAT rules

In this section, we will walk through the steps of setting up HP-UX IPFilter that translate the source IP addresses of all packets from the compute private subnet to the IP address of the gateway node. For addir more sophisticated NAT rules, please refer to the IPFilter documentation.

1. Create a file with NAT rules.

Example 1: Map packets from all Compute Nodes in the 192.168.0.x subnet to a single IP address 15.99.84.23

```
% cat << EOF > /tmp/nat.rules

# lan0 -- interface to the external network
# NAT IP interface: 15.99.84.23
map lan0 192.168.0.0/24 -> 15.99.84.23/32 portmap tcp/udp
40000:60000
map lan0 192.168.0.0/24 -> 15.99.84.23/32
EOF

% cat /tmp/nat.rules

# lan0 -- interface to the external network
# NAT IP interface: 15.99.84.23
map lan0 192.168.0.0/24 -> 15.99.84.23/32 portmap tcp/udp
40000:60000
map lan0 192.168.0.0/24 -> 15.99.84.23/32
```

Example 2: Map packets from specific Compute Nodes (192.168.0.3 and 192.168.0.4) to a single IP address 15.99.84.23

```
% cat << EOF > /tmp/nat.rules

# lan0 -- interface to the external network
# NAT IP interface: 15.99.84.23
map lan0 192.168.0.3/32 -> 15.99.84.23/32 portmap tcp/udp
40000:60000
map lan0 192.168.0.3/32 -> 15.99.84.23/32
map lan0 192.168.0.4/32 -> 15.99.84.23/32 portmap tcp/udp
40000:60000
map lan0 192.168.0.4/32 -> 15.99.84.23/32
EOF

% cat /tmp/nat.rules

# lan0 -- interface to the external network
# NAT IP interface: 15.99.84.23
map lan0 192.168.0.3/32 -> 15.99.84.23/32 portmap tcp/udp
40000:60000
map lan0 192.168.0.3/32 -> 15.99.84.23/32
```

```
map lan0 192.168.0.4/32 -> 15.99.84.23/32 portmap tcp/udp  
40000:60000  
map lan0 192.168.0.4/32 -> 15.99.84.23/32  
EOF
```

More examples of NAT and other IPFilter rules are available at /opt/ipf/examples.

2. Enable NAT based on this rule set

```
% ipnat -f /tmp/nat.rules
```

Note: If there are existing NAT rules that you want to replace, you must flush and delete that rule set before loading the new rules:

```
% ipnat -FC -f /tmp/nat.rules
```

For more complicated manipulations of the rules, refer to ipnat man pages.

Enable NAT in the Compute Nodes

Configure the routing table in each Compute Node to route packets that cannot be directly delivered to the Management Server, which serves as the gateway node. This will normally be done automatically by compute_config.

Example: In this example, lan1 is the private subnet of the Compute Nodes, and the Management Server' lan1 interface is 192.168.0.1. The following steps should be performed to configure the routing tables in each Compute Node:

- On each Compute Node, issue the command:

```
% /usr/sbin/route add default 192.168.0.1 1
```

- On each Compute Node, add or modify the default route in /etc/rc.config.d/netconf. For example:

```
ROUTE_DESTINATION[0]=default  
ROUTE_MASK[0]=""  
ROUTE_GATEWAY[0]="192.168.0.1"  
ROUTE_COUNT[0]=""  
ROUTE_ARGS[0]=""
```

Verify the NAT installation

To verify that the HP-UX IPFilter NAT configuration works, select any Compute Node that has its address translated, and perform the following tests:

- Ping test

On the Compute Node, contact a machine that is not part of the cluster using the 'ping' command. For example:

```
% ping 15.99.22.42
```

If there is no packet loss, then NAT is enabled.

- DISPLAY Server Interaction Test

1. On the Compute Node, set the DISPLAY variable to a display server that is not part of the cluster, for instance your local desktop.

```
% setenv DISPLAY 15.99.22.42:0.0 (if it is csh)
```

2. Try to bring up an xterm on the DISPLAY server:

```
% xterm &
```

If the xterm is brought up in the DISPLAY server, then NAT is enabled.

References:

- [3.6.1 Introduction to NAT \(Network Address Translation\)](#)

[Back to Top](#)

1.3.2 External /home File Server

During the installation, manager_config presents the option to mount a /home directory to all of the nodes in the cluster. The default is to use /home on the Management Server as the home directory for the cluster. If an alternate mount point is used, it is necessary to perform the following steps before starting the "Invoke /opt/clusterpack/bin/manager_config on Management Server" step:

- If it is not already setup, configure the file server to export the directory you intend to mount as /home.
- Connect the file server to the ProCurve 5308xl switch. The file server's connection to the switch should use the same subnet mask as the rest of the Compute Nodes. This may require adding and/or configuring an additional network card on the file server.
- When manager_config prompts you for the mount location of /home, use the IP address of the file server that is accessible to the ProCurve 5308xl switch.

References:

- [Step 8 Copy the License Files to the Management Server](#)

[Back to Top](#)

1.3.3 Adding Head Nodes to an ClusterPack cluster

The default use model of an ClusterPack cluster is that end users will submit jobs remotely through the ClusterWare GUI or by using the ClusterWare CLI from the Management Node. Cluster administrators generally discourage users from logging into the Compute Nodes directly. Users are encouraged to use the Management Server for accessing files and performing routine tasks. When it is desirable to add additional nodes for this purpose, or for more intense computational tasks such as job pre or post processing and compilation, additional "head nodes" can be used. In this document, the term "head node" refers to such user-accessible nodes that allow for interactive use. Head nodes can be included in an ClusterPack Cluster using the following approach:

- The head nodes should include an additional network card to allow the node to be accessible to the wider area network.
- Head nodes should be added to the cluster using the same approach as Compute Nodes. They can be included in the initial cluster definition or added at a later time using the '-a' option to manager_config and compute_config.
- Administrators may choose to close these nodes from running ClusterWare jobs or only make them accessible only to particular queues. (See ClusterWare documentation for more information).
- It may be convenient to use the clgroup command to create groups to represent the head node(s) and the remaining Compute Nodes.
- Use compute_config to configure the additional network cards to allow the head node(s) to be accessible outside of the cluster. Assign the available network cards publicly accessible IP addresses as appropriate to your local networking configuration.

[Back to Top](#)

1.3.4 Set up TCP-CONTROL

ClusterPack delivers a package to allow some control of TCP services coming into the Compute Nodes. This package, called TCP-CONTROL, can be used to limit users from accessing the Compute Nodes directly, but should be used with great care due to several restrictions. TCP-CONTROL can be used to force users to submit jobs through ClusterWare Pro™ only. It accomplishes this by disabling telnet and remsh access to the Compute Nodes from the manager. However, this will also cause several important telnet- and remsh-based applications to fail for non-root users. The tools affected are the multi-system aware tools (clsh, clps, etc.) and the AppRS utilities (apprs_ls, apprs_clean, etc.).

Note:

Enabling TCP-CONTROL by setting the /etc/hosts.deny file will prevent users' access to multi-system aware tools and AppRS utilities.

By default, the TCP-CONTROL package is installed on the Compute Nodes, but is not configured to restrict access in any way. TCP control is restricted by the settings in /etc/hosts.allow and /etc/hosts.deny files on each Compute Node. The /etc/hosts.deny file is initially configured with no entries, but has two commented lines that can be uncommented to prevent users from accessing the Compute Nodes:

```
ALL:ALL@<Management Server name>
```

```
ALL:ALL@<Management Server private IP address>
```

By uncommenting these lines, all users from the Management Server will be denied access. There is also a /etc/hosts.allow file that explicitly permits access to some users. It is configured, by default, to allow access to root and lsfadmin:

```
ALL:root@ALL
```

```
ALL:lsfadmin@ALL
```

Although the hosts.deny file disallows all access, the entries in hosts.allow override the settings of hosts.deny. The hosts.deny file also does not prevent users from accessing telnet and remsh between Compute Nodes. This allows MPI-based applications to run when submitted to a ClusterWare Pro™ queue.

More information about the settings in hosts.deny and hosts.allow can be found in the man pages:

```
% man tcpd
```

```
% man hosts_access
```

[Back to Top](#)

ClusterPack



Software Upgrades and Reinstalls

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

[1.4.1 Software Upgrades and Reinstalls Overview](#)

[1.4.2 Prerequisites for Software Upgrades and Reinstalls](#)

[1.4.3 Reinstallation and Configuration Steps](#)

[1.4.4 Upgrading from Base Edition to Clusterware Edition](#)

[Reinstall Step 1 Obtain New license key\(s\)](#)

[Reinstall Step 2 Invoke /opt/clusterpack/bin/manager_config on Management Server](#)

[Reinstall Step 3 Invoke /opt/clusterpack/bin/compute_config on Management Server](#)

[Reinstall Step 4 Invoke /opt/clusterpack/bin/finalize_config on Management Server](#)

[1.4.5 Upgrading from V2.2 to V2.3](#)

1.4.1 Software Upgrades and Reinstalls Overview

Overview

It is very important to read this entire section before beginning the upgrade or reinstallation process.

As with the installation, ClusterPack uses a three-stage process for reinstalling and configuring an ClusterPack managed cluster.

- Installation and configuration of the Management Server
- Installation and configuration of the Compute Nodes
- Verification of the Management Server and Compute Nodes

In the steps below, it is assumed that the nodes are in a booted state and the MP cards have been configured. It is possible to run mp_register at any time to make changes to the MP interfaces of the nodes currently registered with manager_config. clbootnodes can also be used at any time to ensure that nodes with registered MPs are in a fully booted state.

The Reinstall/Upgrade steps suggested in this section cannot be used to modify the cluster LAN interface addresses or hostnames of the Management Server or Compute Nodes. Changes of this nature can only be accomplished by a complete re-configuration of the cluster (See Initial

Installation and Setup). The reinstallation path is only meant to ensure that all of the ClusterPack software is correctly installed and the cluster layout described by earlier invocations of manager_config is configured correctly.

References:

- [1.2.1 Comprehensive Installation Overview](#)

ClusterPack V2.3 supports an upgrade path from ClusterPack V2.2

[Back to Top](#)

1.4.2 Prerequisites for Software Upgrades and Reinstalls

Prerequisites

The key components of the HP-UX Technical Clusters are:

- Management Server: HP Integrity server with HP-UX 11i Version 2.0 TCOE
- Compute Nodes: HP Integrity servers with HP-UX 11i Version 2.0 TCOE
- Cluster Management Software: ClusterPack V2.3

The following prerequisites are assumed:

- HP-UX 11i v2.0 TCOE is installed on the Management Server
- HP-UX 11i v2.0 TCOE is installed on each Compute Node
- HP-UX 11i Ignite-UX on the Management Server
- HP-UX 11i Ignite-UX on each Compute Node
- Host names and IP addresses are assigned to all Compute Nodes
- All Compute Nodes are reachable from the Management Server

[Back to Top](#)

1.4.3 Reinstallation and Configuration Steps

The steps in this section have to be followed in the specified order to ensure that everything works correctly. The finalize_config tools will be run at the end to validate the installation and configuration steps.

[Back to Top](#)

1.4.4 Upgrading from Base Edition to Clusterware Edition

Upgrading from Base Edition to Clusterware Edition is done using the "forced reinstall" path that is documented below. During `manager_config` you will be given an opportunity to provide a valid Clusterware License key. If you have a key, Clusterware will be installed and integrated into the remaining ClusterPack tools.

Please obtain your Clusterware license key BEFORE reinstalling the ClusterPack software.

[Back to Top](#)

Reinstall Step 1 Obtain New license key(s)

You will be given the option to use the existing license file(s) for the reinstalled cluster. If you wish to retain the current license file(s), nothing more is required.

Upgrading from Base Edition to Clusterware Edition

If you are upgrading from Base Edition to Clusterware Edition, you will need to redeem your Clusterware Edition license certificate using the instructions in 1.2.3 Pre-Install Checklist. You can reuse the ClusterPack license file, and specify a location for the Clusterware license file.

Increasing the size of an existing cluster

If you are performing the forced reinstallation as part of adding more compute nodes to your cluster, you will need to obtain new license files using the instruction on your new license certificates.

During `manager_config`, please do not reuse the existing license files. Instead provide the location of the new license files, and they will be automatically installed into the correct locations.

References:

-
-

[Back to Top](#)

Reinstall Step 2 Invoke `/opt/clusterpack/bin/manager_config` on Management Server

This tool is the main installation and configuration driver. Invoke this tool with "force install" option -F:

```
% /opt/clusterpack/bin/manager_config -F
```

Note:

manager_config will ask for the same software depot that was used the last time the cluster was installed. If you are using the ClusterPack V2.3 DVD as the source, please mount the DVD and have it accessible to the Management server BEFORE invoking manager_config -F

References:

- [Step 3 Allocate File System Space](#)
- [Step 4 Obtain a License File](#)

[Back to Top](#)

Reinstall Step 3 Invoke /opt/clusterpack/bin/compute_config on Management Server

This tool is the driver that installs and configures appropriate components on every Compute Node. It is invoked with the "force install" option -F as follows:

```
% /opt/clusterpack/bin/compute_config -F
```

[Back to Top](#)

Reinstall Step 4 Invoke /opt/clusterpack/bin/finalize_config on Management Server

Finalize and validate the installation and configuration of the ClusterPack software.

```
% /opt/clusterpack/bin/finalize_config
```

This step completes the installation and configuration process, does verification checks on the Cluster Management Software, and validates the installation. It prints out diagnostic error messages if the installation is not successful.

[Back to Top](#)

1.4.5 Upgrading from V2.2 to V2.3

ClusterPack V2.3 supports an upgrade path from ClusterPack V2.2. Customers that currently deploy ClusterPack V2.2 on HP Integrity servers use HP-UX 11i Version 2.0 TCOE. ClusterPack V2.3 provides a mechanism for the use of the majority of V2.2 configuration settings for the V2.3 configuration.

Before starting the upgrade, it is important to have all of your Compute Nodes in good working order. All Compute Nodes and MP cards should be accessible. The LSF queues (if in use) should be empty of all jobs, and the nodes should be idle.

Instructions for upgrading from V2.2 to V2.3:

- Backup the cluster user-level data.
- Install the V2.3 backup utilities.

```
% swinstall -s <depot_with_V2.3> CPACK-BACKUP
```

- Take a backup of the cluster information.

```
% /opt/clusterpack/bin/clbackup -f <backup_file_name>
```

- Copy the backup file to another system for safe keeping.
- Remove the TCP wrappers on your Compute Nodes.

```
% clsh /usr/bin/perl -p -i -e "'s^ /usr/lbin/tcpd^';'" /etc/inetd.conf
```

- Remove the Compute Nodes from the SIM database.

```
% /opt/sysinvmgr/bin/simdevice -r  
' /opt/sysinvmgr/bin/simdevice -l | egrep ^Name: |  
awk '{print "-n", $2}' | grep \.'
```

- Install the new ClusterPack manager software.

```
% swinstall -s <depot_with_V2.3> CPACK-MGR
```

- Run manager_config in upgrade mode using the file you created in Step 3.

```
% /opt/clusterpack/bin/manager_config -u <backup_file_name>
```

- Register your MP cards. (To save time, check out the new -f option to compute_config.)

```
% /opt/clusterpack/bin/mp_register
```

- Install the new software on the Compute Nodes. (The -u is important.)

```
% /opt/clusterpack/bin/compute_config -u
```

- Verify that everything is working as expected.

```
% /opt/clusterpack/bin/finalize_config
```

[Back to Top](#)

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

Copyright 1994-2004 hewlett-packard company

ClusterPack



Golden Image Tasks

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

- [1.5.1 Create a Golden Image of a Compute Node from the Management Server](#)
- [1.5.2 Distribute Golden Image to a set of Compute Nodes](#)

1.5.1 Create a Golden Image of a Compute Node from the Management Server

A system image is an archive of a computer's file system. Capturing the file system of a computer captures the basic state of a computer system. An image does not generally include all files however. By default, /tmp and other temporary files, network directories and host specific configuration files are not included.

A system image may be referred to as a golden image or a recovery image. The different names used to refer to the image reflect the different reasons for creating it. Administrators may create a "recovery" image of a node in the event that the node experiences hardware failure or the file system is accidentally removed or corrupted. Administrators may also create a "golden" image for the purpose of installing it on other nodes to insure that each node in their cluster is configured exactly the way they want.

The HP Technical Cluster includes a tool to simplify the process of creating images:

```
% /opt/clusterpack/bin/sysimage_create
```

The recommended procedure for creating a golden image of a Compute Node is:

- Log into the Compute Node to be archived.
- Perform general file system cleanup/maintenance. For example, it may be desirable to search for and remove core files.

Next, from the Management Server:

- Ensure that the system is not being used. It is advisable that the system stop accepting new LSF jobs while the archive is being made:

```
% badmin hclose <hostname>
```

- In addition, you should either wait until all running jobs complete, or suspend them:

```
% bstop -a -u all -m <hostname>
```

- Execute sysimage_create on the Management Server and pass the name of the file from which you would like the image to be made. For example:

```
% /opt/clusterpack/bin/sysimage_create <hostname>
```

- Monitor the output for possible error conditions. The image will be stored in /var/opt/ignite/archives/<hostname>

After the image creation is complete, the Compute Node should be opened for accepting Clusterware jobs.:

```
% badmin hopen <hostname>
```

If the system you want to archive was not installed as part of the HP Technical Cluster, it will not contain the sysimage_create tool. In this case, the make_sys_command should be used directly to create the golden image. Verify that Ignite-UX is installed on the system:

```
% swlist -l product Ignite-UX
```

If it is not, you will need to obtain and install this product first.

<http://software.hp.com>

Read the man pages for make_sys_image(1m) to find out more about creating system images.

The user can control what files are included in an image through the use of the -l, -g and -f arguments to make_sys_image. See the man pages for make_sys_image(1m) for more information. By default, files in /tmp and most files in /var/adm and /var/tmp are excluded. More importantly, many of the system configuration files in /etc are replaced by "newconfig" versions. The reasoning behind this is that the archive may be used on a different system where the settings from the "golden" machine will be incorrect. A complete list of files which will be excluded from an image can be viewed by using the command:

```
% /opt/ignite/data/scripts/make_sys_image -x -s local
```

Users may wish to exclude additional files by using the -f option to make_sys_image. If the archive is only intended to reinstall the machine where it was made, make_sys_image -l may be used to include all configuration files in their current state.

[Back to Top](#)

1.5.2 Distribute Golden Image to a set of Compute Nodes

To distribute a golden image to a set of Compute Nodes, you need to first register the image. To register the image, use the command:

```
% /opt/clusterpack/bin/sysimage_register <full path of image>
```

If the image was created with sysimage_create, the full path of the image was displayed by sysimage_create. Images are stored in the directory:

```
/var/opt/ignite/archives/<hostname>
```

To distribute the golden image to the Compute Nodes, use the command:

```
% /opt/clusterpack/bin/sysimage_distribute <full path of image>
[ hostname | all ]
```

The keyword "all" can be used to distribute the image to all of the Compute Nodes in the cluster, or a single hostname can be specified. sysimage_distribute will reboot each Compute Node for installation with the specified image.

If the image was sent to a node that was already part of the cluster, that node must have the Compute Node software reconfigured. For more information see the "Software Upgrades and Reinstalls" section.

```
% compute_config -a <node name>
```

If the image was sent to a node that will be added to the cluster, please see the "Add Node(s) to the Cluster" use case.

References:

- [1.4.1 Software Upgrades and Reinstalls Overview](#)
- [1.6.1 Add Node\(s\) to the Cluster](#)

[Back to Top](#)

ClusterPack



System Maintenance Tasks

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

[1.6.1 Add Node\(s\) to the Cluster](#)

[1.6.2 Remove Node\(s\) from the Cluster](#)

[1.6.3 Install Software in Compute Nodes](#)

[1.6.4 Remove Software from Compute Nodes](#)

[1.6.5 Update Software in Compute Nodes](#)

[1.6.6 Add Users to Compute Nodes](#)

[1.6.7 Remove Users from Compute Nodes](#)

[1.6.8 Change System Parameters in Compute Nodes](#)

[1.6.9 Define Compute Node Inventories for Consistency checks](#)

[1.6.10 Define Consistency Check Timetables on Compute Node Inventories](#)

[1.6.11 Compare the Inventories of a Set of Nodes](#)

[1.6.12 Execute remote commands on one or more nodes](#)

[1.6.13 Copy files within nodes in a cluster](#)

[1.6.14 List a user's process status on one or more cluster nodes](#)

[1.6.15 Kill a user's process \(or all of the user's processes\) on some/all Cluster Nodes](#)

[1.6.16 Create a Cluster Group](#)

[1.6.17 Remove a Cluster Group](#)

[1.6.18 Add Nodes to a Cluster Group](#)

[1.6.19 Remove Nodes from a Cluster Group](#)

[1.6.20 Add File Systems to Compute Nodes](#)

[1.6.21 Remove File systems from Compute Nodes](#)

[1.6.22 How is the ClusterPack license server managed?](#)

1.6.1 Add Node(s) to the Cluster

Compute Nodes are added to ClusterPack by using the three stage process that is used for installation.

- Installation and configuration of the Management Server
- Installation and configuration of the Compute Nodes
- Verification of the Management Server and Compute Nodes

The steps in this section have to be followed in the specified order to ensure that everything works correctly.

Step 1 Invoke /opt/clusterpack/bin/manager_config on Management Server

Invoke /opt/clusterpack/bin/manager_config with a "add node" option -a.

```
% /opt/clusterpack/bin/manager_config -a  
<new_node_name>:<new_node_ip_addr>
```

This command adds the new node with the specified hostname and IP address to the cluster. It also reconfigures all of the components of ClusterPack to accommodate the new addition.

The '-a' option can include multiple host:ip pairs if more than one node needs to be added to the system.

For more information on the use of manager_config, refer to the man pages.

```
% man manager_config
```

Step 2 Invoke mp_register on Management Server

If the host being added to the cluster has an MP interface, it should be registered and possibly configured with mp_register:

```
% /opt/clusterpack/bin/mp_register -a <new_node_name>
```

The -a option can be repeated when adding multiple hosts at one time.

The mp_register utility will prompt you for information to configure and/or register an MP card for the host being added.

Step 3 Invoke clbootnodes on Management Server

Boot the node. To assist in booting the node, you may use clbootnodes. If the node is already booted, ensure that it has the IP address you indicated when running manager_config. Nodes can be booted from a "first boot" state or an image of another compute node can be installed on any new node(s). To boot a single node which has HP-UX installed on it, use:

```
% clbootnodes <new_node_name>
```

To boot multiple nodes, use:

```
% clbootnodes
```

In the later case, the utility will prompt you (for each node in the cluster) whether to boot it or skip it.

To boot a compute node with a system image, use the "-i" option to clbootnodes and specify the image. The image must have been created by sysimage_create and registered with

sysimage_register. You can see a list of registered images by executing:

The full path of the image must be given to clbootnodes:

Note:

After installing an image with clbootnodes -i, it may be necessary to wait several minutes after clbootnodes returns before running compute_config as the deamons may need time to start and stablize

Step 4 Invoke /opt/clusterpack/bin/compute_config on Management Server

This tool is the driver that installs and configures appropriate components on every Compute Node. It is invoked with the "add node" option -a as follows:

```
% /opt/clusterpack/bin/compute_config -a <new_node_name>
```

This command configures the new node with the specified hostname to serve as a Compute Node in the cluster.

The '-a' option can be repeated if more than one node needs to be added to the system.

For more information on the usage of compute_config, refer to the man pages.

```
% man compute_config
```

Step 5 Invoke /opt/clusterpack/bin/finalize_config on Management Server

Finalize and validate the addition of the new node to ClusterPack:

```
% /opt/clusterpack/bin/finalize_config
```

This step completes the configuration process, does verification checks on the Cluster Management Software, and validates the addition of the new node.

[Back to Top](#)

1.6.2 Remove Node(s) from the Cluster

Compute Nodes are removed from ClusterPack by using the three stage process that is used for installation:

- Installation and configuration of the Management Server
- Installation and configuration of the Compute Nodes
- Verification of the Management Server and Compute Nodes

The steps in this section must be followed in the specified order to ensure that everything works

correctly.

Step 1 Invoke /opt/clusterpack/bin/manager_config on Management Server

Invoke /opt/clusterpack/bin/manager_config with a "remove node" option -r.

```
% /opt/clusterpack/bin/manager_config -r <node_name>
```

This command removes the node with the specified hostname from the cluster. It also reconfigures all of the components of ClusterPack to accommodate the removal of the node.

The '-r' option can be repeated if more than one node needs to be removed from the system.

For more information on the usage of manager_config, refer to the man pages.

```
% man manager_config
```

Step 2 Invoke /opt/clusterpack/bin/compute_config on Management Server

This tool is the driver that installs and configures appropriate components on every Compute Node. It is invoked with the "remove node" option -r as follows:

```
% /opt/clusterpack/bin/compute_config -r <node_name>
```

The '-r' option can be repeated if more than one node needs to be removed from the system.

For more information on the usage of compute_config, refer to the man pages.

```
% man compute_config
```

Step 3 Invoke /opt/clusterpack/bin/finalize_config on Management Server

Finalize and validate the removal of the node to ClusterPack:

```
% /opt/clusterpack/bin/finalize_config
```

This step does verification checks on the Cluster Management Software, and validates the removal of the node.

[Back to Top](#)

1.6.3 Install Software in Compute Nodes

Using the SCM GUI

To add additional software to Compute Nodes using SCM GUI, do the following:

- Under "Tools", select "Software Management", and then double-click on "Install Software".
- Select the node(s) and/or node group to install on.
- This will bring up the swinstall GUI, from which you can specify the software source and select the software to be installed.

References:

- [3.2.3 How to Run SCM Web-based GUI](#)

Using CLI

Software can also be installed on Compute Nodes using the /opt/clusterpack/bin/clsh tool to run the swinstall command. However, this may not work in a guarded cluster.

- To install product PROD1 on all Compute Nodes

```
% /opt/clusterpack/bin/clsh /usr/sbin/swinstall -s
<depot> PROD1
```

- To install product PROD1 on just the Compute Node group "cae"

```
% /opt/clusterpack/bin/clsh -C
cae /usr/sbin/swinstall -s <depot> PROD1
```

[Back to Top](#)

1.6.4 Remove Software from Compute Nodes

Using the SCM GUI

To add additional software to Compute Nodes using SCM GUI, do the following:

- Under "Tools", select "Software Management", and then double-click on "Uninstall Software".
- Select the node(s) and/or node group to install on.
- This will bring up the swremove GUI, from which you can specify the software source and select the software to be removed.

References:

- [3.2.3 How to Run SCM Web-based GUI](#)

Using the CLI

Software can also be removed from Compute Nodes using the /opt/clusterpack/bin/clsh tool to run the swremove command:

- To remove product PROD1 on all Compute Nodes

```
% /opt/clusterpack/bin/clsh /usr/sbin/swremove PROD1
```

- To install product PROD1 on just the Compute Node group "cae"

```
% /opt/clusterpack/bin/clsh -C cae /usr/sbin/remove  
PROD1
```

[Back to Top](#)

1.6.5 Update Software in Compute Nodes

The process for updating software is the same as for installing software. (See "Install Software in Compute Nodes"). swinstall will verify that the software you are installing is a newer version than what is already present. For patches, and software in non-depot format, it will be necessary to follow the specific directions given with the patch/update.

References:

- [1.6.3 Install Software in Compute Nodes](#)

[Back to Top](#)

1.6.6 Add Users to Compute Nodes

Using the SCM GUI

To add users to the cluster, do the following:

- Select the Management Server.
- Under "Tools", select "System Administration", and then click on "Accounts for Users and Groups".
- This will bring up the user account GUI, where you can specify the user account parameters to use in creating the account.

All user accounts are administered from the Management Server. Any changes to a user's account will be pushed to all the Compute Nodes using NIS.

References:

- [3.2.3 How to Run SCM Web-based GUI](#)

Using the CLI

To add users to the Compute Nodes, first add the user to the Management Server with the useradd command. (man useradd(1M) for more information).

```
% useradd
```

Use ypmake to push the new user's account information to the Compute Nodes:

```
% /var/yp/ypmake
```

[Back to Top](#)

1.6.7 Remove Users from Compute Nodes

Using the SCM GUI

To remove users from the cluster, do the following:

- Select the Management Server.
- Under "Tools", select "System Administration", and then click on "Accounts for Users and Groups".
- This will bring up the user account GUI, where you can specify the user account to remove.

All user accounts are administered from the Management Server. Any changes to a users account will be pushed to all the Compute Nodes using NIS.

References:

- [3.2.3 How to Run SCM Web-based GUI](#)

Using the CLI

User accounts should be removed from the Management Server as normal with userdel (man userdel(1M) for more information).

```
% userdel
```

After removing the user, use ypmake to push this change to the Compute Nodes.

```
% /var/yp/ypmake
```

[Back to Top](#)

1.6.8 Change System Parameters in Compute Nodes

Using the SCM GUI:

- Select one or more nodes.
- Under "Tools", select "System Administration", and then click on "System Properties".
- A SAM System Properties window will appear for each node selected.

For greater efficiency and consistency, perform this operation only on a single Compute Node, and then a golden image be created from that Compute Node and pushed to the other Compute Nodes.

References:

- [3.2.3 How to Run SCM Web-based GUI](#)
- [1.5.1 Create a Golden Image of a Compute Node from the Management Server](#)
- [1.5.2 Distribute Golden Image to a set of Compute Nodes](#)

[Back to Top](#)

1.6.9 Define Compute Node Inventories for Consistency checks

To define Compute Node inventories for consistency checks use the SCM GUI to access the SIM GUI.

Using the SCM GUI:

- Select one or more nodes.
- Under "Tools", select "System Inventory", and then click "SysInvMgr portal".
- This launches the SIM GUI.

References:

- [3.2.3 How to Run SCM Web-based GUI](#)

Using the SIM GUI:

- Log in as "admin".
- Select the "Filter" folder.
- Click on "Create Filter".
- Select a name to uniquely identify the inventory filter.
- Enter an optional description.
- Select one or more categories (e.g. System, Memory, I/O Devices).
- Select one or more Groups from the selected categories (e.g. BundleContents,

LogicalMemory).

- For fine control over inventory collection, use "Advanced Settings" to select or unselect specific items.

[Back to Top](#)

1.6.10 Define Consistency Check Timetables on Compute Node Inventories

To define Compute Node inventories for consistency checks use the SCM GUI to access the SIM GUI.

Using the SCM GUI:

- Select one or more nodes.
- Under "Tools", select "System Inventory", and then click "SysInvMgr portal".
- This launches the SIM GUI.

References:

- [3.2.3 How to Run SCM Web-based GUI](#)

Using the SIM GUI:

- Log in as "admin".
- Select the "Schedule" folder.
- Click on "Create Schedule".
- Select a name to uniquely identify the inventory schedule.
- Enter an optional description.
- Enter a start date.
- Enter an end date.
- Enter a collection frequency.

[Back to Top](#)

1.6.11 Compare the Inventories of a Set of Nodes

To define Compute Node inventories for consistency checks use the SCM GUI to access the SIM GUI.

Using the SCM GUI:

- Select one or more nodes.
- Under "Tools", select "System Inventory", and then click "SysInvMgr portal".

- This launches the SIM GUI.

References:

- [3.2.3 How to Run SCM Web-based GUI](#)

Using the SIM GUI:

- Log in as "admin".
- Select the "Filter" folder.
- Click "Create Filter".
- Enter a name to uniquely identify the inventory filter.
- Enter an optional description.
- Select one or more categories (e.g. System, Memory, I/O Devices).
- Select one or more Groups from the selected categories (e.g. BundleContents, LogicalMemory).
- For fine control over inventory collection, use "Advanced Settings" to select or unselect specific items.

[Back to Top](#)

1.6.12 Execute remote commands on one or more nodes

A remote command can be executed on one or more nodes in the cluster from any node by using the 'clsh' command in /opt/clusterpack/bin.

Some examples of clsh usage are:

- Invoke 'uname -a' on all cluster nodes

```
% clsh uname -a
```

- Invoke 'uname -a' on node1

```
% clsh -C node1 uname -a
```

- Invoke 'uname -a' on all nodes in the cluster group 'sub1'

```
% clsh -C sub1 uname -a
```

- Invoke 'uname -a' on node1 and node3

```
% clsh -C node1+node3 uname -a
```

For more details on the usage of clsh, invoke the command.

```
% man clsh
```

[Back to Top](#)

1.6.13 Copy files within nodes in a cluster

The 'clcp' command in /opt/clusterpack/bin is used to copy files between cluster nodes. Each file or directory argument is either a remote file name of the form "%h:path" or "cluster:path" or a local file name (containing no ':' characters).

Some examples of clcp usage are:

- Update /etc/checklist on all nodes with the local /etc/checklist

```
% clcp /etc/checklist %h:/etc/checklist
% clcp /etc/checklist cluster:/etc/checklist
```

- Update /etc/checklist on node1 and node3 with the local /etc/checklist

```
% clcp -C node1+node3 /etc/checklist %
h:/etc/checklist
```

- Copy multiple local files to all nodes

```
% clcp a.txt b.txt c.txt %h:/tmp/.
```

- Copy multiple remote files to multiple local files

```
% clcp %h:/tmp/a.txt /tmp/a.%h.txt
```

For more details on the usage of clcp, invoke:

```
% man clcp
```

[Back to Top](#)

1.6.14 List a user's process status on one or more cluster nodes

The 'clps' command in /opt/clusterpack/bin is used to produce a 'ps' output that includes the host name. A clps command with no arguments lists all the processes associated with the user invoking the command on all Compute Nodes.

Some examples of clps usage are:

- List all processes belonging to user 'joeuser'

```
% clps -u joeuser
```

- List all processes on node3 and node4

```
% clps -C node3+node4 -a
```

For more details on the usage of clps, invoke the command:

```
% man clps
```

[Back to Top](#)

1.6.15 Kill a user's process (or all of the user's processes) on some/all Cluster Nodes

The 'clkill' command in /opt/clusterpack/bin is used to kill processes on Cluster Nodes. Since using PIDs on a cluster is not feasible given there will be different PIDs on different hosts, clkill can kill processes by name.

Some examples of clps usage:

- Kill all processes belonging to user 'joeuser'

```
% clkill -u joeuser
```

- Interactively kill all processes named 'view_server'

```
% clkill -i -r view_server
will result in a session like:
node0 2260 user1 ? 0:00 view_server
Kill (y/n)? n
node1 8070 user1 ? 0:04 view_server
Kill (y/n)? y
```

- Kill a process with PID 2260 on node1

```
% clkill -C node1 -p 2260
```

For more details on the usage of clkill, invoke the command:

[Back to Top](#)

1.6.16 Create a Cluster Group

Groups of Compute Nodes can be created, and added to all tools in ClusterPack using /opt/clusterpack/bin/clgroup.

The following example creates a node group "cae" containing compute cluster nodes "lucky000", "lucky001", and "lucky002":

```
% /opt/clusterpack/bin/clgroup -a cae lucky000 lucky001 lucky002
```

clgroup can also form groups from existing groups. For more details on the usage of clgroup, invoke the command:

```
% man clgroup
```

[Back to Top](#)

1.6.17 Remove a Cluster Group

Groups of Compute Nodes can be removed from ClusterPack using /opt/clusterpack/bin/clgroup.

The following example removes the node group "cae":

```
% /opt/clusterpack/bin/clgroup -r cae
```

Note that the above-mentioned command just removes the group; the nodes are still part of the cluster, and users can submit jobs to the nodes.

For more details on the usage of clgroup, invoke the command:

```
% man clgroup
```

[Back to Top](#)

1.6.18 Add Nodes to a Cluster Group

Compute Nodes can be added to existing groups in ClusterPack using /opt/clusterpack/bin/clgroup.

The following example adds nodes "lucky006" and "lucky008" to the node group "cae" :

```
% /opt/clusterpack/bin/clgroup -a cae lucky006 lucky008
```

Groups can also be created or extended using the name of a pre-existing group. For more details on the usage of clgroup, invoke the command:

```
% man clgroup
```

[Back to Top](#)

1.6.19 Remove Nodes from a Cluster Group

Compute Nodes can be removed from existing groups in ClusterPack using /opt/clusterpack/bin/clgroup.

The following example removes node "lucky006" from the node group "cae" :

```
% /opt/clusterpack/bin/clgroup -r cae lucky006
```

Groups can also have entire groups of nodes removed by using the name of a pre-existing group. For more details on the usage of clgroup, invoke the command:

```
% man clgroup
```

[Back to Top](#)

1.6.20 Add File Systems to Compute Nodes

The file system for Compute Nodes can be defined using System Administration Manager (SAM). Invoke SAM from the command line or from within SCM and select "Disks and File Systems". Invoke SAM from the command line or from within SCM and select "Disks and File Systems". Select "Actions->Add Local File System->Using the Logical Volume Manager" and enter the required information. Repeat this operation for each Compute Node.

[Back to Top](#)

1.6.21 Remove File systems from Compute Nodes

Removing file systems can be done in a similar fashion as adding file systems (See "Add File Systems to Compute Nodes"). From SAM, select the file system you want to remove and select "Actions->Remove". Do this for each node in the cluster.

References:

- [1.6.20 Add File Systems to Compute Nodes](#)

[Back to Top](#)

1.6.22 How is the the ClusterPack license server managed?

ClusterPack Base Edition

The ClusterPack Base Edition license server is based on FlexLM licensing technology.

The Base Edition license server is installed and configured by the manager_config tool. The license server is started by manager_config, and it is installed to start during a normal system boot.

To manually start the ClusterPack license server:

```
% /sbin/init.d/cpack.server start
```

To manually stop the ClusterPack license server:

```
% /sbin/init.d/cpack.server stop
```

ClusterPack Clusterware Edition

Every installation of ClusterPack Clusterware Edition includes a fully functional Base Edition license manager. All Base Edition license server functions should be used to manage that portion of the license server.

Platform Computing's Clusterware Pro V5.1 uses a proprietary licensing scheme. For more information on managing the Clusterware Pro license functionality, Please see the "Platform Computing Clusterware Pro V5.1 Overview"

References:

- [3.7.5 How do I start and stop the Clusterware Pro V5.1 daemons?](#)

[Back to Top](#)

ClusterPack



System Monitoring Tasks

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

[1.7.1 Get an Overview of Cluster Health](#)

[1.7.2 Get an Overview of the Job Queue Status](#)

[1.7.3 Get details on health of specific Compute Nodes](#)

[1.7.4 View Usage of Resources in Compute Node\(s\)](#)

[1.7.5 Monitor Compute Nodes based on resource thresholds](#)

1.7.1 Get an Overview of Cluster Health

Using finalize_config

The best way to determine the overall health of the cluster is to run:

```
% finalize_config
```

Finalize_config performs a series of tests to determine the overall health of the individual components of the cluster that have been automatically setup and administered by ClusterPack. Finalize_config can be run repeatedly without side effects.

The health of the cluster for accepting and running jobs can also be determined using tools provided as part of Clusterware Pro.

Using the Clusterware Pro V5.1 Web Interface:

The Hosts Tab, on the left hand side of the screen, contains a table showing information about your hosts' resources. The "Detailed View" shows the current State and Batch State. The Detailed View is accessed by selecting:

```
% View->Details
```

There are two different indications of health and status: State, and Batch State.

- State refers to the state of the host.
- Batch State refers to the state of the host, and the state of the daemons running on that host. A detailed list of batch states is shown below.

For more information, select the online help:

- Select Help->Platform Help
- Select "View" under the "Hosts" section in the left hand pane.
- Select "Change your hostview" to see a description of the icons.

Using the Clusterware Pro V5.1 CLI:

Status of each node is available using the CLI on the Management Server:

```
% bhosts
```

STATUS shows the current status of the host and the SBD daemon. Batch jobs can only be dispatched to hosts with an ok status. A more detailed list of STATUS is available in the long report:

```
% bhosts -1
```

or

```
% bhosts -1 <hostname>
```

The lsload command provides an instantaneous view of the load state of the Compute Nodes.

```
% lsload
```

A more detailed list of the load information is available in the long report:

```
% lsload -1
```

or

```
% lsload -1 <hostname>
```

Common Terms

Both the Web interface and the CLI use the same terms for the health and status of the individual Compute Nodes. These terms are used to define the Batch State of an individual node.

- ok - The host is up
- closed_Adm - The host is not accepting new jobs. It has been closed by the administrator
- closed_Lock - The host is not accepting jobs and all jobs running on it have been suspended. It has been locked by the administrator.
- closed_Busy - The host is not accepting new jobs. Some load indices

- have exceeded their thresholds.
- closed_Excl - The host is not accepting jobs until the exclusive job running on it completes.
- closed_Full - The host is not accepting new jobs. The configured maximum number of jobs that can run on it has been reached.
- closed_Wind - The host is not accepting jobs. The dispatch window that has been defined for it is closed.
- unlicensed - The host is not accepting jobs. It does not have a valid LSF license for sbatchd and LIM is down.
- closed_Lim - The LIM on the host is unreachable, but sbatchd is running.
- unavail - The LIM on the host is not responding and sbatchd is down.
- unreach - The sbatchd is down.
- unreach - The sbatchd is down and LIM is unreachable

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

1.7.2 Get an Overview of the Job Queue Status

Using the Clusterware Pro V5.1 Web Interface:

Select the Queues tab. An overview of available job queues is displayed. The following details are displayed:

- State - The state of the queue. Any queue with an "Open:Active" state can accept new jobs.
- Jobs Pending - The number of jobs pending in that queue.
- Jobs Running - The number of jobs running in that queue.
- Priority - The relative priority of the queue.

Using the Clusterware Pro V5.1 CLI:

From the CLI issue the following command.

% bqueues

By default bqueues returns the following information about all queues: queue name, queue priority, queue status, job slot statistics, and job state statistics.

Additional information is available for all hosts in the cluster, or for a single hostname.

% bqueues -1

or

```
% bqueues -l <hostname>
```

For more information, see the man page:

```
% man bqueues
```

Common Terms

Both the Web interface and the CLI use the same terms for the health and status of the job submission queues. These terms are used to define the State of an individual queue.

- Open - The queue is able to accept jobs.
- Closed - The queue is not able to accept jobs.
- Active - Jobs in the queue may be started.
- Inactive - Jobs in the queue cannot be started for the time being.

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

1.7.3 Get details on health of specific Compute Nodes

Using the Clusterware Pro V5.1 Web Interface:

The Hosts Tab, located on the left hand side of the screen, contains a table showing information about your hosts' resources. The "Detailed View" shows the current Stage and Batch State. The Detailed View is accessed by selecting:

```
% View->Details
```

There are two different indications of health and status: State, and Batch State.

- State refers to the state of the host.
- Batch State refers to the state of the host, and the state of the daemons running on that host. A detailed list of batch states is shown below.

For more information, select the online help:

```
Help->Platform Help
```

Using the Clusterware Pro V5.1 CLI:

Default status from each node is available using:

```
% bhosts <hostname>
```

STATUS shows the current status of the host and the SBD daemon. Batch jobs can only be dispatched to hosts with an ok status. A more detailed list of STATUS is available in the long report:

```
% bhosts -l <hostname>
```

Detailed information for the specific host can be received from most of the CLI tools:

```
% lshosts -l <hostname>
% lsload -l <hostname>
```

Common Terms

Both the Web interface and the CLI use the same terms for the health and status of the individual Compute Nodes. These terms are used to define the Batch State of an individual node.

- ok - The host is up
- closed_Adm - The host is not accepting new jobs. It has been closed by the administrator
- closed_Lock - The host is not accepting jobs and all jobs running on it have been suspended. It has been locked by the administrator.
- closed_Busy - The host is not accepting new jobs. Some load indices have exceeded their thresholds.
- closed_Excl - The host is not accepting jobs until the exclusive job running on it completes.
- closed_Full - The host is not accepting new jobs. The configured maximum number of jobs that can run on it has been reached.
- closed_Wind - The host is not accepting jobs. The dispatch window that has been defined for it is closed.
- unlicensed - The host is not accepting jobs. It does not have a valid LSF license for sbatchd and LIM is down.
- closed_Lim - The LIM on the host is unreachable, but sbatchd is running.
- unavail - The LIM on the host is not responding and sbatchd is down.
- unreach - The sbatchd is down.

References:

- [1.7.1 Get an Overview of Cluster Health](#)
- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

1.7.4 View Usage of Resources in Compute Node(s)

Using the Clusterware Pro V5.1 Web Interface:

From the Hosts Tab:

- Select the host to be monitored using the checkbox next to each host. More than one host can be selected.
- From the menu select Host->Monitor
- A new window will open that displays the current resource usage of one of the selected hosts. Four resources are displayed: total system memory, CPU Utilization, swap space available, and /tmp space available.
- If more than one host was selected, you can change between hosts using the Monitor Selected Hosts menu.
- The individual resources are samples every 60 seconds. It may be several minutes before the graphs show any information for a given node.

Using the Clusterware Pro V5.1 CLI:

Resources available for job scheduling can be seen using the following command:

```
% bhosts
```

This will display a report for all the Compute Nodes in the cluster. To get the resource usage for an individual Compute Node, specify the name of the node on the command line:

```
% bhosts -1 <hostname>
```

For more information, see the man page:

```
% man bhosts
```

Static resources of the node can be seen using the following command:

```
% lshosts
```

This will display a report for all the Compute Nodes in the cluster. To get the static resource information for an individual Compute Node, specify the name of the node on the command line:

```
% lshosts -1 <hostname>
```

For more information, see the man page:

```
% man lshosts
```

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)

- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

1.7.5 Monitor Compute Nodes based on resource thresholds

Using the Clusterware Pro V5.1 Web Interface:

From the Hosts Tab

- From the View menu select View->Choose Columns
- Add the Available Column resource to the Displayed Columns list.
- Click OK
- The new resource to be monitored will be displayed on the Host tab screen.

Using the Clusterware Pro V5.1 CLI:

Using the lshosts command, a resource can be specified. Only hosts that meet the resource requirement will be displayed.

```
% lshosts [-R res_req] [hostname]
```

For example, to find all the hosts with at least 4096MB of available memory:

```
% lshosts -R mem=4096
```

Membership in logical groups defined with the clgroup command can also be given as a resource.

```
% lshosts -R <group_name>
```

For a full list of currently defined resources, use the lsinfo command:

```
% lsinfo
```

See the lshosts man page for more information.

```
% man lshosts
```

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

ClusterPack



Workload Management Tasks

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

- [1.8.1 Add new Job Submission Queues](#)
- [1.8.2 Remove Queues](#)
- [1.8.3 Restrict user access to specific queues](#)
- [1.8.4 Add resource constraints to specified queues](#)
- [1.8.5 Change priority of specified queues](#)
- [1.8.6 Add pre/post run scripts to specified queues](#)
- [1.8.7 Kill a job in a queue](#)
- [1.8.8 Kill all jobs owned by a user](#)
- [1.8.9 Kill all jobs in a queue](#)
- [1.8.10 Suspend a job in a queue](#)
- [1.8.11 Suspend all jobs owned by a user](#)
- [1.8.12 Suspend all jobs in a queue](#)
- [1.8.13 Resume a suspended job in a queue](#)
- [1.8.14 Resume all suspended jobs owned by a user](#)
- [1.8.15 Resume all suspended jobs in a queue](#)

1.8.1 Add new Job Submission Queues

A new queue can be added to the cluster by editing the file: /share/platform/clusterware/conf/lsbatch/<clustername>/configdir/lsb.queues

The name of your cluster can be determined by using the Clusterware Pro V5.1 CLI:

```
% lsid
```

This above file should contain several initial queues that can be copied and modified as desired. The top of the lsb.queues file contains comments about the fields of a queue description.

If the queue you are adding is for AppRS, several fields in the queue description must be carefully set to the values required by AppRS. An example AppRS queue definition can be found in /opt/apprs/examples/example_queue.

After adding, removing or modifying queues, it is necessary to reconfigure LSF to read the new queue information. This is done from the Management Server using the Clusterware Pro V5.1 CLI:

```
% badmin reconfig
```

Verify the queue has been added by using the Clusterware Pro V5.1 CLI:

```
% bqueues -l <queue_name>
```

References:

- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

1.8.2 Remove Queues

A queue can be removed from the cluster by editing the file: /share/platform/clusterware/conf/lsbatch/<clustername>/configdir/lsb.queues

The name of your cluster can be determined by using the Clusterware Pro V5.1 CLI:

```
% lsid
```

Before removing a queue, it should be closed using the Clusterware Pro V5.1 CLI:

```
% badmin qclose <queue name>
```

Jobs still executing can be killed or allowed to run to completion before removing the queue.

Delete or comment out the queue definitions that you want to remove.

After adding, removing or modifying queues, it is necessary to reconfigure LSF to read the new queue information. This is done from the Management Server using the Clusterware Pro V5.1 CLI:

```
% badmin reconfig
```

Verify the queue has been removed by using the Clusterware Pro V5.1 CLI:

```
% bqueues -l <queue_name>
```

References:

- [1.8.1 Add new Job Submission Queues](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

1.8.3 Restrict user access to specific queues

Using the Clusterware Pro V5.1 CLI:

The file /share/platform/clusterware/conf/lsbatch/<clustername>/configdir/lsb.queues controls which users can submit to a specific queue.

The name of your cluster can be determined by using the Clusterware Pro V5.1 CLI:

```
% lsid
```

Edit the lsb.queues file and look for a USERS line for the queue you wish to restrict. If a USERS line exists, you can add or remove users from it. Otherwise, add a line of the form:

```
USERS = <keyword all or a list of userid's>
```

to the queue definition.

After adding, removing or modifying queues, it is necessary to reconfigure LSF to read the new queue information. This is done from the Management Server using the Clusterware Pro V5.1 CLI:

```
% badmin reconfig
```

Verify the queue has been modified by using the Clusterware Pro V5.1 CLI:

```
% bqueues -l <queue_name>
```

References:

- [1.8.1 Add new Job Submission Queues](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

1.8.4 Add resource constraints to specified queues

To add resource constraints to a queue, edit the file /share/platform/clusterware/conf/lsbatch/<clustername>/configdir/lsb.queues.

The name of your cluster can be determined by using the Clusterware Pro V5.1 CLI:

```
% lsid
```

Find the queue definition you wish to modify. The following entries for maximum resource usage can be modified or added for each queue definition:

- CPULIMIT = minutes on a host
- FILELIMIT = file size limit
- MEMLIMIT = bytes per job
- DATALIMIT = bytes for data segment
- STACKLIMIT = bytes for stack
- CORELIMIT = bytes for core files
- PROCLIMIT = processes per job

RES_REQ is a resource requirement string specifying the condition for dispatching a job to a host. Resource reservation and locality can also be specified in this string.

Groups of Compute Nodes created using clgroup are treated as resource requirements in Clusterware Pro. If a group name is added as a RES_REQ to a queue, jobs submitted to that queue will only run on nodes that are members of that group.

After adding, removing or modifying queues, it is necessary to reconfigure LSF to read the new queue information. This is done from the Management Server using the Clusterware Pro V5.1 CLI:

```
% badmin reconfig
```

Verify the queue has been modified by using the Clusterware Pro V5.1 CLI:

```
% bqueues -l <queue_name>
```

References:

- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

1.8.5 Change priority of specified queues

The file /share/platform/clusterware/conf/lsbatch/<clustername>/configdir/lsb.queues controls the priority of each queue.

The name of your cluster can be determined by using the Clusterware Pro V5.1 CLI:

```
% lsid
```

Edit the lsb.queues file and look for a PRIORITY line for the queue you wish to change. If a PRIORITY line exists, you can change the value specified. Otherwise, add a line of the form:

```
PRIORITY = <integer value>
```

to the queue definition. Queues with higher priority values are searched first during scheduling.

After adding, removing or modifying queues, it is necessary to reconfigure LSF to read the new queue information. This is done from the Management Server using the Clusterware Pro V5.1 CLI:

```
% badmin reconfig
```

Verify the queue has been modified by using the Clusterware Pro V5.1 CLI:

```
% bqueues -l <queue_name>
```

References:

- [1.8.1 Add new Job Submission Queues](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

1.8.6 Add pre/post run scripts to specified queues

The file /share/platform/clusterware/conf/lsbatch/<clustername>/configdir/lsb.queues controls the pre and post commands associated with each queue.

The name of your cluster can be determined by using the Clusterware Pro V5.1 CLI:

```
% lsid
```

Pre-execution commands are executed before a job is run from the queue. Post-execution commands are executed when a job successfully completes execution from the queue. This can be useful for acquiring and releasing special resources such as access to hardware devices or software licenses.

Note:

Application Restart already makes use of these pre and post execution commands. Therefore, AppRS queues cannot have their PRE_EXEC and POST_EXEC values modified.

Edit the lsb.queues file and look for PRE_EXEC or POST_EXEC lines for the queue you wish to change. If such lines exist, you can change the values specified. Otherwise, add lines of the form:

```
PRE_EXEC = <command or full path to a script>
POST_EXEC = <command or full path to a script>
```

to the queue definition. The command or tool should be accessible and runnable on all nodes that the queue services.

After adding, removing or modifying queues, it is necessary to reconfigure LSF to read the new queue information. This is done from the Management Server using the Clusterware Pro V5.1 CLI:

```
% badmin reconfig
```

Verify the queue has been modified by using the Clusterware Pro V5.1 CLI:

```
% bqueues -l <queue_name>
```

References:

- [1.8.1 Add new Job Submission Queues](#)

[Back to Top](#)

1.8.7 Kill a job in a queue

Using the Clusterware Pro V5.1 CLI:

Jobs can be killed using the bkill command:

```
% bkill <jobid>
```

Users can kill their own jobs. Queue administrators can kill jobs associated with a particular queue.

References:

- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

1.8.8 Kill all jobs owned by a user

Using the Clusterware Pro V5.1 CLI:

All of a user's jobs can be killed by using the bkill on the special job id 0:

```
% bkill -u <username> 0
```

Users can kill their own jobs. Queue administrators can kill jobs associated with a particular queue.

References:

- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

1.8.9 Kill all jobs in a queue

Using the Clusterware Pro V5.1 CLI:

All of the jobs in a queue can be killed by using the bkill command with the -q option:

```
% bkill -q <queue name> -u all 0
```

Users can kill their own jobs. Queue administrators can kill jobs associated with a particular queue.

References:

- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

1.8.10 Suspend a job in a queue

Using the Clusterware Pro V5.1 CLI:

```
% bstop <jobid>
```

Users can suspend their own jobs. Queue administrators can suspend jobs associated with a particular queue.

References:

- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

1.8.11 Suspend all jobs owned by a user

Using the Clusterware Pro V5.1 CLI:

All of a user's jobs can be suspended using the special 0 job id:

```
% bstop -u <userid> 0
```

Users can suspend their own jobs. Queue administrators can suspend jobs associated with a particular queue.

References:

- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

1.8.12 Suspend all jobs in a queue

Using the Clusterware Pro V5.1 CLI:

All of the jobs in a queue can be suspended by a queue administrator using the special 0 job id:

```
% bstop -q <queue name> -u all 0
```

References:

- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

1.8.13 Resume a suspended job in a queue

Using the Clusterware Pro V5.1 CLI:

```
% bresume <jobid>
```

Users can resume their own jobs. Queue administrators can resume jobs associated with a particular queue.

References:

- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

1.8.14 Resume all suspended jobs owned by a user

Using the Clusterware Pro V5.1 CLI:

All of a user's jobs can be resumed using the Clusterware Pro V5.1 CLI by using the special 0 job id:

```
% bresume -u <userid> 0
```

Users can resume their own jobs. Queue administrators can resume jobs associated with a particular queue.

References:

- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

1.8.15 Resume all suspended jobs in a queue

Using the Clusterware Pro V5.1 CLI:

All of the jobs in a queue can be resumed by a queue administrator using the special 0 job id:

```
% bresume -q <queue name> -u all 0
```

References:

- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

ClusterPack



System Troubleshooting Tasks

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

- [1.9.1 Locate a Compute Node that is down](#)
- [1.9.2 Get to the console of a Compute Node that is down](#)
- [1.9.3 Bring up a Compute Node with a recovery image](#)
- [1.9.4 View system logs for cause of a crash](#)
- [1.9.5 Bring up the Management Server from a crash](#)
- [1.9.6 Troubleshoot SCM problems](#)
- [1.9.7 Replace a Compute Node that has failed with a new machine](#)

1.9.1 Locate a Compute Node that is down

Any node that shows a state of "unavail" or "unreach" is potentially down and should be checked by a system administrator. In order to determine the state of nodes on the cluster, the tools should be used.

Using the Clusterware Pro V5.1 Web Interface

The default hosts view is a table showing information about your hosts' resources. The default view is accessed from:

View->Details

There are two different indications of health and status: State, and Batch State.

- State refers to the state of the host.
- Batch State refers to the state of the host, and the state of the daemons running on that host.

Using the Clusterware Pro V5.1 CLI

An overview of a nodes health can be obtained by looking at the Clusterware Pro status and batch states of that node.

```
% lshosts -l <hostname>
% bhosts -l <hostname>
```

References:

- [1.7.1 Get an Overview of Cluster Health](#)
- [1.7.3 Get details on health of specific Compute Nodes](#)
- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

1.9.2 Get to the console of a Compute Node that is down

If a Compute Node is unreachable using the Management Server LAN within the cluster, it may be necessary to use the MP card to access the console of the node.

- Telnet to the remote MP-console of the Compute Node

```
% telnet <node1-mp>
```

- Use the MP interface to view any diagnostic messages from the Compute Node.

References:

- [3.7.1 What is Clusterware Pro?](#)

[Back to Top](#)

1.9.3 Bring up a Compute Node with a recovery image

Recovery images created with /opt/clusterpack/bin/sysimage_create are stored in /var/opt/ignite/archives/<hostname>, where <hostname> is the name of the node from which the image was taken. The images are stored in files based on the time in which the image was created. To distribute such an image to a Compute Node, use the tool /opt/clusterpack/bin/sysimage_register to register the image with Ignite-UX:

```
% /opt/clusterpack/bin/sysimage_register <full path of image>
```

To bring up a Compute Node using this image, you should use the sysimage_distribute script:

```
% /opt/clusterpack/bin/sysimage_distribute <full path of image>
<hostname>
```

This will reboot the machine, hostname, and cause the machine to install from the golden image you specified.

References:

- [1.5.2 Distribute Golden Image to a set of Compute Nodes](#)

[Back to Top](#)

1.9.4 View system logs for cause of a crash

The system logs are located in /var/admin/syslog/syslog.log

The crash logs are stored in /var/adm/crash

The installation and configuration logs for ClusterPack are stored in /var/opt/clusterpack/log

[Back to Top](#)

1.9.5 Bring up the Management Server from a crash

After a crash, the Management Server state can be checked by running:

```
% /opt/clusterpack/bin/finalize_config
```

[Back to Top](#)

1.9.6 Troubleshoot SCM problems

There are two common problems that are discussed here. For any additional troubleshooting help, please see:

- Planning, installing, and updating ServiceControl Manager

<http://docs.hp.com/hpux/onlinedocs/5187-1882/5187-1882.html>

- ServiceControl Manager Technical Reference

<http://docs.hp.com/hpux/pdf/B8339-90030.pdf>

When I try to add a node, I get "Properties file for <xyz> doesn't exist."

Solution:

- Make sure that the hostname is fully qualified in /etc/hosts on both the Management Server and the managed node, if it exists in /etc/hosts, and that any shortened host names are aliases instead of primary names. For example:
 - 10.1.2.3 cluster.abc.com cluster
- should be used instead of:
 - 10.1.2.3 cluster
- Make sure that AgentConfig is installed on the managed node, and that mxrmi and mxagent are running.
- ps -ef | grep mx should produce something like this:

```
root 23332 1 0 15:42:17 ? 1:08 /opt/mx/lbin/mxagent
root 23334 1 0 15:42:17 ? 0:59 /opt/mx/lbin/mxrmi
root 24269 24252 1 01:30:51 pts/0 0:00 grep mx
```
- If AgentConfig is installed and running, uninstall it and then reinstall it:

```
% /usr/sbin/swremove AgentConfig
```

- To install AgentConfig, type;

```
% /usr/sbin/swinstall -s <CMS>:/var/opt/mx/depot11
AgentConfig
```

- where <CMS> is the hostname of the Management Server.

scmgr prints "out of memory" errors.

Solution:

- On the Management Server, using SAM or kmtune, make sure that the Kernel Configurable Parameter max_thread_proc is at least 256, and that nkthread is at least 1000.

[Back to Top](#)

1.9.7 Replace a Compute Node that has failed with a new machine

If a Compute Node fails due to a hardware problem, and must be replaced, the new node can be added to the cluster using the IP address and hostname of the failed node or can be added

with a new name and IP address.

Replacing with a new hostname and IP address

In this case, the replacement node is handled simply by removing the failed node and adding the new node. Remove the failed node from the cluster using the following commands:

```
% manager_config -r <bad_node_name>
% compute_config -r <bad_node_name>
```

The nodes MP will automatically be removed from the MP register database. Add the replacement node to the cluster using the same approach as adding a new node to the cluster:

```
% manager_config -a <new_node_name>:<new_node_ip_addr>
% mp_register -a <new_node_name>
% clbootnodes <new_node_name>
% compute_config -a <new_node_name>
% Add the new node into groups as appropriate using clgroups
```

Replacing with the same hostname and IP address

If the hostname and IP Address from the failed node will be assigned to the replacement node do NOT remove the failed node from the cluster using the '-r' option. This will remove the node from any groups that have been setup, and it will remove any automated SIM information collections that refer to that node.

The replacement node should be installed, and connected to the cluster. Use mp_register to configure the MP card on the new host and/or register the MP card's IP address:

```
% mp_register -a <new_node_name>
```

Boot the node using clbootnodes

```
% clbootnodes <new_node_name>
```

Finally, run compute_config to update this new node:

```
% compute_config -a <new_node_name>
```

References:

-
-

[Back to Top](#)

ClusterPack



Job Management Tasks

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

[2.1.1 Invoke the Workload Management Interface from the Management Server](#)

[2.1.2 Invoke the Workload Management Interface from the intranet](#)

[2.1.3 Prepare for job submission](#)

[2.1.4 Submit a job to a queue](#)

[2.1.5 Submit a job to a group](#)

[2.1.6 Set a priority for a submitted job](#)

[2.1.7 Check the status of a submitted job](#)

[2.1.8 Check the status of all submitted jobs](#)

[2.1.9 Examine data files during a job run](#)

[2.1.10 Register for notification on completion of a submitted job](#)

[2.1.11 Kill a submitted job in a queue](#)

[2.1.12 Kill all jobs submitted by the user](#)

[2.1.13 Kill all jobs submitted by the user in a queue](#)

[2.1.14 Suspend a submitted job in a queue](#)

[2.1.15 Suspend all jobs submitted by the user](#)

[2.1.16 Suspend all jobs submitted by the user in a queue](#)

[2.1.17 Resume a suspended job in a queue](#)

[2.1.18 Resume all suspended jobs submitted by the user](#)

[2.1.19 Resume all suspended jobs submitted by the user in a queue](#)

[2.1.20 Submit a MPI job in a queue](#)

[2.1.21 Suspend a submitted MPI job](#)

[2.1.22 Resume a suspended MPI job](#)

2.1.1 Invoke the Workload Management Interface from the Management Server

Using the Clusterware Pro V5.1 Web Interface:

- Set the DISPLAY environment variable to your display. For example, using csh or tcsh:

```
% setenv DISPLAY mypc:0
```

- Go to the following URL in the web browser:

```
% /opt/netscape/netscape
http://<management_server>:8080/Platform/login/Login.jsp
```

- Enter your Unix user name and password.

This assumes that the gaadmin services have been started by the LSF Administrator.

Note:

The user submitting a job must have access to the Management Server and to all the Compute Nodes that will execute the job. To prevent security problems, the super user account (i.e. root) cannot submit any jobs using .

References:

- [3.7.6 How do I start and stop the Clusterware Pro V5.1 Web GUI?](#)
- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)

[Back to Top](#)

2.1.2 Invoke the Workload Management Interface from the intranet

Using the Clusterware Pro V5.1 Web Interface:

- Go to the following URL in a web browser:

```
%  
http://<management_server>:8080/Platform/login/Login.jsp
```

- Enter your Unix user name and password.

This assumes that the gaadmin services have been started by the LSF Administrator.

References:

- [3.7.6 How do I start and stop the Clusterware Pro V5.1 Web GUI?](#)
- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)

[Back to Top](#)

2.1.3 Prepare for job submission

Using the Clusterware Pro V5.1 Web Interface:

From the jobs tab:

- Select Job->Submit.
- Enter job data.
- Click Submit.

Data files required for the job may be specified using the '-f' option to the bsub command. This optional information can be supplied on the "Advanced" tab within the Job Submission screen. For an explanation of the '-f' options please see "Transfer a file from intranet to specific Compute Nodes in the cluster".

Using the Clusterware Pro V5.1 CLI:

```
% bsub <options> command <arguments>
```

Data files may be specified using the '-f' option to bsub. See "Transfer a file from intranet to specific Compute Nodes in the cluster" for more information on using the '-f' option to transfer files within the cluster.

Jobs may be submitted to a Group of Compute Nodes (if the group was created using the clgroup tool) by specifying a resource requirement of the group name:

```
% bsub -R <group_name> command <arguments>
```

See bsub(1) man page for complete syntax.

```
% man 1 bsub
```

References:

- [2.2.3 Transfer a file from intranet to specific Compute Nodes in the cluster](#)
- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

2.1.4 Submit a job to a queue

Using the Clusterware Pro V5.1 Web Interface:

From the Jobs tab:

- Select Job->Submit.
- Enter relevant Job information.
- Select Queue name from Queue list.

Using the Clusterware Pro V5.1 CLI:

```
% bsub -q <queue_name> <cmd>
```

Use bqueues to list available Queues.

```
% bqueues
```

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

2.1.5 Submit a job to a group

Using the Clusterware Pro V5.1 Web Interface:

From the Jobs tab:

- Select Job->Submit.
- Enter relevant Job information.
- Select the "Resources" tab. Enter the group name in the "Resource Requirement string" field.

Using the Clusterware Pro V5.1 CLI:

```
% bsub -R <group_name> <cmd>
```

Use clinfo to list the current groups, and their membership.

```
% clinfo
```

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

2.1.6 Set a priority for a submitted job

Using the Clusterware Pro V5.1 Web Interface:

Set a priority at submission by:

- From the Jobs Tab
- Select Job->Submit.
- Using the Queue pull down menu, select a queue with a high priority.

After submission:

- From the Jobs Tab.
- Select a job from the current list of pending jobs.
- Select Job->Switch Queue.
- Switch the job to a queue with a higher priority

The relative priority of the different Queues can be found on the "Queue Tab".

Using the Clusterware Pro V5.1 CLI:

```
% bmod -sp <priority> <job_ID>
% bswitch <desintationqueue> <job_ID>
```

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

2.1.7 Check the status of a submitted job

Using the Clusterware Pro V5.1 Web Interface:

From the Jobs tab:

- Select Tools->Find.
- On the Find dialog enter:
 - Select ID from the Field list.
 - Select = from the Condition list.
 - Enter the Job ID in the Value field.
- Click Find.

Using the Clusterware Pro V5.1 CLI:

```
% bjobs <job_ID>
% bjobs -l <job_ID>
```

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

2.1.8 Check the status of all submitted jobs

Using the Clusterware Pro V5.1 Web Interface:

From the Jobs tab:

- Review the Jobs table.
- Use the Previous and Next buttons to view more jobs.

Using the Clusterware Pro V5.1 CLI:

```
% bjobs
% bjobs -l
```

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

2.1.9 Examine data files during a job run

Using the Clusterware Pro V5.1 CLI:

```
% bpeek <job_ID>
```

References:

- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

2.1.10 Register for notification on completion of a submitted job

Using the Clusterware Pro V5.1 Web Interface:

From the Jobs tab:

- Select Job->Submit.
- Click Advanced.
- Select done from Send email notification when job is
- Enter the email address in the email to field.

Using the Clusterware Pro V5.1 CLI:

Using the CLI, users are automatically notified when a job completes.

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

2.1.11 Kill a submitted job in a queue

Using the Clusterware Pro V5.1 Web Interface:

From the Jobs tab:

- Select the job from the Jobs table.
- Select Jobs->Kill.

Using the Clusterware Pro V5.1 CLI:

`% bkill <job_ID>`

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

2.1.12 Kill all jobs submitted by the user

Using the Clusterware Pro V5.1 Web Interface:

From the Jobs tab:

- Select Tools->Find.
- Select User from the Field list.
- Type the user name in the Value field.
- Click Find.
- Click Select All.
- Click Kill.

Using the Clusterware Pro V5.1 CLI:

```
% bkill -u <username> 0
```

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

2.1.13 Kill all jobs submitted by the user in a queue

Using the Clusterware Pro V5.1 Web Interface:

From the Jobs tab:

- Select Tools->Find.
- Select the Advanced tab.
- Select User from the Field list in the Define Criteria section.
- Type the user name in the Value field.
- Click << to add to the list.
- Select Queue from the Field list.
- Select the queue from the Queues list.
- Click Add to List.
- Click Find.
- Click Select All.
- Click Kill.

Using the Clusterware Pro V5.1 CLI:

```
% bkill -q <queue name> -u <username> 0
```

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

2.1.14 Suspend a submitted job in a queue

Using the Clusterware Pro V5.1 Web Interface:

From the Jobs tab:

- Select the job from the Jobs table.
- Select Job->Suspend.

Using the Clusterware Pro V5.1 CLI:

```
% bstop <job_ID>
```

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

2.1.15 Suspend all jobs submitted by the user

Using the Clusterware Pro V5.1 Web Interface:

From the Jobs tab:

- Select Tools->Find.
- Select User from the Field list.
- Type the user name in the Value field.
- Click Find.
- Click Select All.
- Click Suspend.

Using the Clusterware Pro V5.1 CLI:

```
% bstop -u <username> 0
```

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

2.1.16 Suspend all jobs submitted by the user in a queue

Using the Clusterware Pro V5.1 Web Interface:

From the Jobs tab:

- Select Tools->Find.
- Select the Advanced tab.
- Select User from the Field list in the Define Criteria section.
- Type the user name in the Value field.
- Click <<
- Select Queue from the Field list.
- Select the queue from the Queue list.
- Click <<
- Click Find.
- Click Select All.
- Click Suspend.

Using the Clusterware Pro V5.1 CLI:

```
% bstop -u <username> -q <queuename> 0
```

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

2.1.17 Resume a suspended job in a queue

Using the Clusterware Pro V5.1 Web Interface:

From the Jobs tab:

- Select the suspended job from the Jobs table.

- Select Job->Resume.

Using the Clusterware Pro V5.1 CLI:

```
% bresume <job_ID>
```

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

2.1.18 Resume all suspended jobs submitted by the user

Using the Clusterware Pro V5.1 Web Interface:

From the Jobs tab:

- Select Tools->Find.
- Select the Advanced tab.
- Select User from the Field list in the Define Criteria section.
- Type the user name in the Value field.
- Click <<
- Select State from the Field list.
- Select Suspend from the State list.
- Click <<
- Click Find.
- Click Select All.
- Click Resume.

Using the Clusterware Pro V5.1 CLI:

```
% bresume -u <username> 0
```

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

2.1.19 Resume all suspended jobs submitted by the user in a queue

Using the Clusterware Pro V5.1 Web Interface:

From the Jobs tab:

- Select Tools->Find.
- Select the Advanced tab.
- Select User from the Field list in the Define Criteria section.
- Type the user name in the Value field.
- Click <<
- Select Queue from the Field list.
- Select the queue from the Queue list.
- Click <<
- Click Find.
- Click Select All.
- Click Resume.

Using the Clusterware Pro V5.1 CLI:

```
% bresume -u <username> -q <queuename> 0
```

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

2.1.20 Submit a MPI job in a queue

Using the Clusterware Pro V5.1 Web Interface:

From the Jobs tab:

- Select Job->Submit.
- Enter the number of processors required in the Max Processors field.
- Complete job data.
- Click Submit.

Using the Clusterware Pro V5.1 CLI:

```
% bsub -n <num_proc> command
```

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)

- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

2.1.21 Suspend a submitted MPI job

Using the Clusterware Pro V5.1 Web Interface:

From the Jobs tab:

- Select the job from the Jobs table.
- Select Job->Suspend.

Using the Clusterware Pro V5.1 CLI:

```
% bstop <job_ID>
```

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

2.1.22 Resume a suspended MPI job

Using the Clusterware Pro V5.1 Web Interface:

From the Jobs tab:

- Select the suspended job from the Jobs table.
- Select Job->Resume.

Using the Clusterware Pro V5.1 CLI:

```
% bresume <job_ID>
```

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

ClusterPack



File Transfer Tasks

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

- [2.2.1 Transfer a file from intranet to the Management Server in the cluster](#)
- [2.2.2 Transfer a file from intranet to all Compute Nodes in the cluster](#)
- [2.2.3 Transfer a file from intranet to specific Compute Nodes in the cluster](#)
- [2.2.4 Transfer a file from a Compute Node to a system outside the cluster](#)
- [2.2.5 Transfer a file from a Compute Node to another Compute node in the cluster](#)
- [2.2.6 Transfer a file from a node to a set of nodes in the cluster](#)

2.2.1 Transfer a file from intranet to the Management Server in the cluster

Using the Clusterware Pro V5.1 Web Interface:

By default, all files transferred using the Web interface will be placed in:

```
% /share/platform/clusterware/tomcat/webapps/Clusterware/users/<userid>
```

From the Jobs tab:

- Tools->Upload/Download Files
- Complete information.
- Click Upload.

Using the CLI:

FTP the file to the Management Server. From the Management Server, ftp to a remote host:

```
% ftp <remote host>
```

References:

- [2.2.5 Transfer a file from a Compute Node to another Compute node in the cluster](#)
- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)

[Back to Top](#)

2.2.2 Transfer a file from intranet to all Compute Nodes in the cluster

If the cluster is a Guarded Cluster, this operation is done in two steps:

- FTP the file to the Management Server.
- Copy the file to all nodes in the cluster.

```
% clcp /a/input.data %h:/date/input.data
% clcp /a/input.data cluster:/date/input.data
```

For more details on the usage of clcp, invoke the command:

```
% man clcp
```

References:

- [2.2.1 Transfer a file from intranet to the Management Server in the cluster](#)

[Back to Top](#)

2.2.3 Transfer a file from intranet to specific Compute Nodes in the cluster

If the cluster is a guarded cluster, this operation must be done in two steps:

- First, FTP the file to the Head node (Management Server)
- Second, distribute the file to specific nodes. There are two methods that can be used:
 1. Use Clusterware Pro V5.1 CLI to distribute the file to the specific nodes that need the file

```
% bsub -f <local_file> op <remote_file>
```

Where 'op' is an operator. An operator specifies whether the file is copied to the remote host, or whether it is copied back from the remote host. The operator must be surrounded by white space. The following describes the operators:

> Copies the local file to the remote file before the job starts. Overwrites the remote file if it exists.

```
% bsub -f <local_file> > <remote_file>
```

< Copies the remote file to the local file after the job completes.
Overwrites the local file if it exists.

```
% bsub -f <local_file> < <remote_file>
```

<< Appends the remote file to the local file after the job
completes. The local file must exist.

```
% bsub -f <local_file> << <remote_file>
```

>< Copies the local file to the remote file before the job starts.
Overwrites the remote file if it exists. Then copies the remote
file to the local file after the job completes. Overwrites the local
file.

```
% bsub -f <local_file> >< <remote_file>
```

<> Copies the local file to the remote file before the job starts.
Overwrites the remote file if it exists. Then copies the remote
file to the local file after the job completes. Overwrites the local
file.

```
% bsub -f <local_file> <> <remote_file>
```

2. Copy the file to specific nodes in the cluster using clcp.

```
% clcp -C node1+node3 /a/input.data %
h:/date/input.data
```

For more details on the usage of clcp, invoke the command:

```
% man clcp
```

References:

- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

2.2.4 Transfer a file from a Compute Node to a system outside the cluster

If the cluster is a guarded cluster, this operation is done in two steps:

- Copy the file from the specified node to the Management Server in the cluster
using clcp.

```
% clcp node1:/a/output.data /tmp/output.data
```

- FTP the file from the Head node to the external target.

References:

- [Guarded Cluster](#)

[Back to Top](#)

2.2.5 Transfer a file from a Compute Node to another Compute node in the cluster

The 'clcp' command in /opt/clusterpack/bin is used to copy files between cluster nodes.

This command can be invoked either from the Management Server or any Compute Node.

[From the Management Server]

```
% clcp node1:/a/data node2:/b/data
```

[Back to Top](#)

2.2.6 Transfer a file from a node to a set of nodes in the cluster

The 'clcp' command in /opt/clusterpack/bin is used to copy files between cluster nodes. Each file or directory argument is either a remote file name of the form "%h:path" or "cluster:path" or a local file name (containing no ':' characters).

Some examples of clcp usage are:

- Update /etc/checklist on all nodes with the local /etc/checklist

```
% clcp /etc/checklist %h:/etc/checklist
% clcp /etc/checklist cluster:/etc/checklist
```

- Update /etc/checklist on node1 and node3 with the local /etc/checklist

```
% clcp -C node1+node3 /etc/checklist %h:/etc/checklist
```

- Copy multiple local files to all nodes

```
% clcp a.txt b.txt c.txt %h:/tmp/.
```

- Copy multiple remote files to multiple local files

```
% clcp %h:/tmp/a.txt /tmp/a.%h.txt
```

For more details on the usage of clcp, invoke the command:

```
% man clcp
```

[Back to Top](#)

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

Copyright 1994-2004 hewlett-packard company

ClusterPack



Miscellaneous Tasks

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

- [2.3.1 Run a tool on a set of Compute Nodes](#)
- [2.3.2 Check resource usage on a Compute Node](#)
- [2.3.3 Check Queue status](#)
- [2.3.4 Remove temporary files from Compute Nodes](#)
- [2.3.5 Prepare application for checkpoint restart](#)
- [2.3.6 Restart application from a checkpoint if a Compute Node crashes](#)
- [2.3.7 Determine if the application fails to complete](#)
- [2.3.8 Check impact on the job if a Compute Node crashes](#)
- [2.3.9 Get a high-level view of the status of the Compute Nodes](#)

2.3.1 Run a tool on a set of Compute Nodes

A set of multi-system aware tools has been provided for use on the cluster. To execute a command on multiple hosts, follow the examples below:

- To run a tool on all the Compute Nodes

```
% clsh <script>
```

- To run a tool on host001

```
% clsh -C host001 <script>
```

- To run a tool on host001 and host00n

```
% clsh -C host001+host00n <script>
```

For more information on the clsh command, please see the man page:

```
% man clsh
```

Alternatively, the tools can be used to submit a job to multiple computer nodes.

Using the Clusterware Pro V5.1 Web Interface:

From the Jobs tab:

- Select Jobs->Submit.
- Enter job information.
- Click Advanced.
- On the Advanced dialog, enter script details in the Pre-execution command field.
- Click OK.
- Click Submit.

Using the CLI:

```
% bsub E 'pre_exec_cmd [args ...]' command
```

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

2.3.2 Check resource usage on a Compute Node

Using the Clusterware Pro V5.1 Web Interface:

From the Jobs tab:

- Select the job from the Jobs table.
- Select Jobs->Monitor.
- Review the charts.

Using the Clusterware Pro V5.1 CLI:

```
% lsload [-1] <host_name>
```

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

2.3.3 Check Queue status

Using the Clusterware Pro V5.1 Web Interface:

From the Jobs tab:

- Review the Queues table. Use the Previous and Next buttons to view more Queues.

Using the Clusterware Pro V5.1 CLI:

```
% bqueues [<queue_name>]
```

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

2.3.4 Remove temporary files from Compute Nodes

The process of removing temporary files created by an application can vary depending on whether the job was submitted to an AppRS queue or a non-AppRS LSF queue.

Jobs submitted to an AppRS queue execute in temporary directories allocated by AppRS on each node where the job runs. AppRS copies the contents of these directories back to the host that the job was submitted from when the job completes. As long as the application only generates files within its execution directory, there is no need for the user to remove temporary files generated by an application.

In the event AppRS restarts an application on a new set of nodes, the original working directories and files created before the migration are not removed. This is done in order to be as careful as possible about avoiding data loss. In addition, the files and directories on the unavailable host(s) cannot be removed at that time because the host cannot be reached. The user must manually remove these files using the AppRS commands `apprs_ls` and `apprs_clean`. Use `apprs_ls` to list the directories allocated by AppRS to the current user:

```
% apprs_ls
```

For non-root users, `apprs_clean` can only be used to remove files that belong to a specific job id. This protects files that are currently in use by another job from being accidentally erased.

```
% apprs_clean <jobid>
```

root can delete all `apprs` related directories. This command should be used with great caution

and should not be used while AppRS jobs are running.

```
% apprs_clean all
```

For jobs submitted to non-AppRS queues, the user's job submission script should include commands to remove files that are no longer needed when the job completes. In the event that the job fails to run to completion it may be necessary to remove these files manually. To find out what hosts the job executed on use the command:

```
% bhist -l <jobid>
```

Included in the output is the list of hosts that the job executed on and the working directory used for execution. This information can be used to manually delete files from a job that was unable to complete successfully.

References:

- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

2.3.5 Prepare application for checkpoint restart

Any job submitted to an AppRS enabled queue is restarted on a new set of hosts if:

- Any host allocated to the job becomes unavailable or unreachable by the other hosts while the job is executing.
- The job is explicitly migrated using the LSF command `bmig`.
- The user's job exits with exit code 3.

(For more information on exit values, see the HP Application ReStart User's Guide)

As long as an application can generate restart files and be restarted from those files, AppRS will ensure that files marked as Highly Available are present when the application is restarted. AppRS will requeue any application that exits with a status of either 2 or 3. If the application (or script that invokes the application) should not be requeued, an exit status other than 2 or 3 should be used.

A job submission script for a checkpoint/restart application should follow the example in `/opt/apprs/examples/job_template`:

```
#!/bin/sh
#BSUB -n 2 # Number of processors requested
#BSUB -e test.stderr # Standard error file
#BSUB -o test.stdout # Standard output file
#BSUB -q normal_apprs
#APPRS INPUT # list input files separated by spaces
#APPRS HIGHLYVISIBLE # list HV (progress) files
#APPRS HIGHLYAVAILABLE # list HA (restart) files
```

```

#APPRS TARGETUTIL 1.0
#APPRS TARGETTIME 10
#APPRS REDUNDANCY 4
# Your job goes here:
if [ "$APPRS_RESTART" = "Y" ]; then
# job as it is run under restart conditions
else
# job as it is run under normal conditions
fi

```

The names of all files that need to be present for the application to run from a restart should be listed with the HIGHLYAVAILABLE tag:

```
#APPRS HIGHLYAVAILABLE <list of files>
```

Other AppRS options can be set in the job submission script. For full details, refer to the

[HP Application ReStart User's Guide](#)

Alternatively, the toolset can be used to trigger checkpointing by your application.

Using the Clusterware Pro V5.1 Web Interface:

From the Jobs tab:

- Select Jobs->Submit.
- Enter job information.
- Click Advanced.
- On the Advanced dialog:
 - Select Checkpoint.
 - Specify an checkpoint period in the every [] minutes field.
 - Specify a checkpoint directory in the directory field.
- On the Advanced dialog, enter script details in the Pre-execution command field.
- Click OK.
- Click Submit.

Using the Clusterware Pro V5.1 CLI:

```
% bsub -k "chkpnt_dir [period] [method=name]" command
```

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

2.3.6 Restart application from a checkpoint if a Compute Node crashes

If a Compute Node crashes, jobs submitted to an AppRS queue will automatically be restarted on a new node or set of nodes as those resources become available. No user intervention is necessary.

[Back to Top](#)

2.3.7 Determine if the application fails to complete

The job state of EXIT is assigned to jobs that end abnormally.

Using the Clusterware Pro V5.1 Web Interface:

From the Jobs tab:

- Review the job states in the Jobs table.
- Use the Previous and Next buttons to view more Jobs.

Using the Clusterware Pro V5.1 CLI:

```
% bjobs <job_ID>
```

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

2.3.8 Check impact on the job if a Compute Node crashes

In the event that a Compute Node crashes or becomes unavailable, it may be desirable to check on jobs that may be affected by the situation.

Using the Clusterware Pro V5.1 CLI:

- List your current and recently finished jobs:

```
% bjobs -a
```

- Request information on a particular job:

```
% bhist <jobid>
```

- or for more information:

```
% bhist -l <jobid>
```

- For jobs submitted to an AppRS queue, details of the job, including failover progress can be viewed using the command:

```
% apprs_hist <jobid>
```

References:

- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

2.3.9 Get a high-level view of the status of the Compute Nodes

Using the Clusterware Pro V5.1 Web Interface:

From the Jobs tab:

- Review the Hosts table.
- Use the Previous and Next buttons to view more hosts.

Using the Clusterware Pro V5.1 CLI:

```
% bhosts
```

References:

- [3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)
- [3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[Back to Top](#)

ClusterPack



Cluster Management Utility Zone Overview

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

[3.1.1 What is Cluster Management Utility Zone?](#)

[3.1.2 What are the Easy Install Tools?](#)

[3.1.3 What are the system imaging tools?](#)

[3.1.4 What are the MSA \(Multi System Aware\) Tools?](#)

[3.1.5 clsh - Runs commands on one, some, or all nodes in the cluster.](#)

[3.1.6 clcp - Copies files to one, some, or all cluster nodes.](#)

[3.1.7 cluptime - Works like ruptime, only for all the nodes in the cluster.](#)

[3.1.8 clps - Cluster-wide ps command](#)

[3.1.9 clkkill - Kills specified processes on specified nodes.](#)

[3.1.10 clinfo - Shows nodes and cluster information.](#)

[3.1.11 clgroup - Creates a logical cluster group of nodes.](#)

[3.1.12 clbroadcast - Telnet and MP based broadcast commands on cluster nodes.](#)

3.1.1 What is Cluster Management Utility Zone?

ClusterPack includes several utilities which can aide both in administrative tasks and in workload management. The multi-system aware tools allow users to easily perform such tasks as:

- Run a command on all nodes in the cluster
- See a list of all of the processes on a cluster owned by a user
- Kill all jobs owned by a user on the entire cluster
- Copy a local file to all nodes in the cluster

These are only a sample of the some of many types of tasks that the multi-system aware tools can perform.

[Back to Top](#)

3.1.2 What are the Easy Install Tools?

The ClusterPack suite includes a set of utilities for setting up a cluster of Itanium 2 nodes. The tools manager_config, mp_register, clbootnodes, compute_config and finalize_config are key components for establishing and administering an Itanium 2 cluster. In particular, these utilities provide:

- An easy step-by-step process for establishing a cluster
- Installation and configuration of ClusterPack software
- Configuration of commonly used HP-UX services
- Configuration of optional MP interfaces
- Networking configuration
- Verification of the cluster's installation and configuration
- The ability to add and remove nodes from the cluster

The process of setting up a cluster is divided into three phases. manager_config prepares the Management Server to act as the central point of management for the cluster. The Compute Nodes are installed and configured with compute_config. finalize_config verifies the installation.

In addition, the tools mp_register and clbootnodes can be used to register and configure MP interfaces and then use those interfaces to automate the booting of nodes.

By default, manager_config interactively asks the user for an IP address range to assign to the Compute Nodes. It is also possible to pass a file containing names and IP addresses to manager_config.

The EasyInstall utilities can also be used to add or remove Compute Nodes from a cluster. manager_config and compute_config can be run with -a and -r options to add or remove individual nodes respectively. In the event that installed software becomes corrupt or otherwise unusable, manager_config can be run with a -F option that will only reinstall software.

References:

-

[Back to Top](#)

3.1.3 What are the system imaging tools?

Ignite-UX is a powerful tool for managing software. Ignite-UX includes the capacity to create an image of a system for the purpose of recovery or replication to other systems. During installation and configuration, ClusterPack configures the cluster's Management Server to act as an Ignite-UX server. ClusterPack also includes three scripts to simplify the process of creating and distributing golden images:

- sysimage_create

- sysimage_register
- sysimage_distribute

These scripts use ClusterPack's knowledge of the cluster configuration to simplify the creation and distribution of system (golden) images. With the use of scripts, creating and distributing images is as simple as running these three tools and providing the name of a host and/or path of the image.

References:

- [1.5.1 Create a Golden Image of a Compute Node from the Management Server](#)
- [1.5.2 Distribute Golden Image to a set of Compute Nodes](#)

[Back to Top](#)

3.1.4 What are the MSA (Multi System Aware) Tools?

These utilities enable invoking common tasks on multiple nodes of the cluster from a single point. The following is a brief description of the utilities. For detailed information, refer to the man pages.

- clsh - Runs commands on one, some or all nodes in the cluster.
- clcp - Copies files to one/some/all cluster nodes.
- cluptime - Works like ruptime, only for all nodes in the cluster.
- clps - Cluster-wide ps command.
- clkill - Kills specified processes on specified nodes.
- clinfo - Shows nodes and cluster information.
- clgroup - Create a logical cluster group of nodes and propagate the group information across all of the management tools in the cluster (PCC Clusterware, SCM, SIM, etc.)

[Back to Top](#)

3.1.5 clsh - Runs commands on one, some, or all nodes in the cluster.

The clsh program runs a remote shell to each host in the compute cluster with the specified command. The command to run is echoed to the user's screen each time one is launched. A new command will not begin until the previous one is finished, i.e. these do not run in parallel.

Sending a SIGINT (usually a ^C) will cause the current host to be skipped, and sending a SIGQUIT (usually a ^\) will immediately abort the whole clsh command.

Percent interpolation, as in clcp, is also supported.

clsh exits non-zero if there are problems running the remote shell commands. A summary of hosts on which problems occurred is printed at the end.

clsh is used as follows:

```
% clsh [-C cluster-group] [options] cmd [args]
```

Examples

To grep for something on all hosts in the cluster:

```
% clsh grep pattern files ...
```

To append something to a file on all machines:

```
% clsh -i "cat >> file" < addendum
```

To run a command with a five second timeout on all the hosts in the cluster group "hp", directing output into separate files:

```
% clsh -o -t5 -C hp date
```

```
% clsh -o -t5 hp date
```

A cluster name without a -C must follow all flag arguments.

For more details on the usage of clsh, invoke the command:

```
% man clsh
```

[Back to Top](#)

3.1.6 clcp - Copies files to one, some, or all cluster nodes.

clcp copies files between nodes in the cluster using rcp. Each file or directory argument is either a remote file name of the form "%h:path", "cluster:path", or a local file name (containing no ":" characters).

clcp can do the following types of copies:

- single local to single local

```
% clcp src dst
```

- single local to multiple local

```
% clcp src dst.%h
```

- single local to multiple remote

```
% clcp src dst:%h or clcp src cluster-group:dst
```

- multiple local to multiple remote

```
% clcp src dst.%h %h:dst
```

- multiple remote to multiple local

```
% clcp %h:src dst.%h
```

Examples

1. Assume that the file /etc/checklist needs to be updated on all HP hosts. Also assume that this file is different on all hosts. The following is a way in which this can be done:

```
% clcp %h:/etc/checklist checklist.%h
% vi checklist.*
```

Make necessary changes.

```
% clcp checklist.%h %h:/etc/checklist
```

If the CLUSTER environment variable was defined as host0+host1, then the above would map to:

```
% rcp host0:/etc/checklist checklist.host0
% rcp host1:/etc/checklist checklist.host1
% vi checklist.host0 checklist.host1
% rcp checklist.host0 host1:/etc/checklist
% rcp checklist.host1 host2:/etc/checklist
```

2. Another way to do the same thing using "%c" instead of "%h" is:

```
% clcp %h:/etc/checklist checklist.%c
% vi checklist.*
```

Make necessary changes.

```
% clcp checklist.%c %h:/etc/checklist
```

which maps to:

```
% rcp host0:/etc/checklist checklist.0
```

```
% rcp host1:/etc/checklist checklist.1  
  
% vi checklist.0 checklist.1  
  
% rcp checklist.0 host0:/etc/checklist  
  
% rcp checklist.1 host1:/etc/checklist
```

3. The following is an example if log files are needed:

```
% clcp %h:/usr/spool/mqueue/syslog %h/syslog.%Y%M%  
D.%T
```

This would save the files in directories (which are the host names) with file names of the form: YYMMDD.TT:TT. The above might map to:

```
% rcp host0:/usr/spool/mqueue/syslog  
host0/syslog.921013.14:43  
  
% rcp host1:/usr/spool/mqueue/syslog  
host1/syslog.921013.14:43
```

4. Like `rcp`, `clcp` can copy many files to the cluster. This is done by:

```
% clcp src1 src2 src3 %h:
```

or

```
% clcp src1 src2 src3 cluster-group:
```

For more details on the usage of `clcp`, invoke the command:

```
% man clcp
```

[Back to Top](#)

3.1.7 **cluptime** - Works like **ruptime**, only for all the nodes in the cluster.

The `cluptime` program reports on the status (similar to `uptime`) of all the hosts in the cluster or specified logical cluster.

`cluptime` is used as follows:

```
% cluptime [ [-C] cluster-group]
```

For more details on the usage of `cluptime`, invoke the command:

```
% man cluptime
```

[Back to Top](#)

3.1.8 clps - Cluster-wide ps command

clps and clkill are the same program with clps producing a "ps" output that includes the host name and clkill allowing processes to be killed.

clps is used as follows:

```
% clps [-C] cluster [-ad] {tty user command pid regexp}
```

For more details on the usage of clps, invoke the command:

```
% man clps
```

[Back to Top](#)

3.1.9 clkill - Kills specified processes on specified nodes.

clps and clkill are the same program with clps producing a "ps" output that includes the host name and clkill allowing processes to be killed. Since using PIDs on a cluster is not feasible given there will be different hosts, clkill can kill processes by name. The -i option should be passed to clkill to allow interactive killing; i.e. any matches are shown and the user is prompted if they wish to kill that process.

clkill is used as follows:

```
% clkill [[-C] cluster] [-signal] [-ilnadv] {tty user command pid regexp}
```

For more details on the usage of clkill, invoke the command:

```
% man clkill
```

[Back to Top](#)

3.1.10 clinfo - Shows nodes and cluster information.

The clinfo command lists which hosts make up a cluster. By default, with no arguments, the current cluster is listed. Non-flag arguments are interpreted as cluster names. Three different output modes are supported.

- Short format (enabled by the -s option)

The short format lists the cluster (followed by a colon) and the hosts it contains; one cluster per line. Long lines do not wrap. If there is only one cluster to be listed and the -v option has not been used, the leading cluster and colon are omitted. This is the default mode if the output is not to a tty device; facilitating the use of clinfo as a component in a larger script.

- Medium format (enabled by the -m option)

The medium format is tabular. It is easier to read, but harder on programs. Long lines wrap and the cluster name is always given, even when there is only one cluster. This is the default mode if the output is to a tty device, like the user's screen.

- Long format (enabled by the -l option)

The long format is essentially a dump of the internal database maintained by cladmin. The cluster name is always output, followed by one record per host. Each field of the record occurs by itself on a line, a blank line separates each record.

The -a option instructs clinfo to display information about all defined clusters, regardless of the -C option or the CLUSTER environment variable setting.

The -C is entirely optional, and is supplied only for consistency with other cluster commands.

clinfo is used as follows:

```
% clinfo [-a] [-v] [-s|-m|-l] [[-C] cluster ...]
```

For more details on the usage of clinfo, invoke the command:

```
% man clinfo
```

[Back to Top](#)

3.1.11 clgroup - Creates a logical cluster group of nodes.

The clgroup command allows the user to create, modify, remove, and list node groups in the core tools of ClusterPack; ServiceControl Manager (SCM), PCC ClusterWare Pro™, Systems Inventory Manager (SIM), and the Cluster Management Utility Zone. Node groups are collections of nodes that are subsets of the entire node membership of the compute cluster. They may have overlapping memberships such that a single node may be a member of more than one group. The node grouping mechanism allows flexible partitioning of a compute cluster into logical collections that match their use model.

The first form of this command allows the user to add node groups to a compute cluster. The initial definition of the node group can be specified as a list of individual nodes and/or other groups. When a previously existing group is used in the formation of a new group, all members of the pre-existing group are added to the new group.

The second form allows the user to remove a node group or nodes from a node group. When removing nodes from a group, the nodes to be removed can be specified in terms of a list of individual nodes and/or other groups. When a previously existing group is specified, all members of that group are removed from the group being modified.

The third form allows the information regarding one or more node groups to be provided in a file.

The last form lists all the node groups in the compute cluster.

clgroup is used as follows:

```
% clgroup -a groupname node1|group [node2|group ...]  
% clgroup -r groupname [node1|group node2|group ...]  
% clgroup -l groupname  
% clgroup
```

Examples

The following command creates a node group "group1" containing compute cluster nodes "nodea", "nodeb", and "nodec".

```
% clgroup -a group1 nodea nodeb nodec
```

To modify group "group1" to replace node "nodec" with node "noded", the following command can be used:

```
% clgroup -r group1 nodec  
% clgroup -a group1 noded
```

A new group "group2" can be formed from all of the nodes in "group1" and the additional nodes "nodex" and "nodey" using the command:

```
% clgroup -a group2 group1 nodex nodey
```

The member nodes of group "group1" may be listed by using:

```
% clgroup -l group1
```

For more details on the usage of **clgroup**, invoke the command:

```
% man clgroup
```

3.1.12 clbroadcast - Telnet and MP based broadcast commands on cluster nodes.

The clbroadcast command is used to broadcast commands to various nodes in the cluster using the Management Processor (MP) interface or telnet interface.

The tool opens a window with a telnet or an MP connection on each target and another "console window" with no echo where all input keyboard actions will be broadcast in all target windows. To send a command to a specific target, type directly in the target window and the command is not broadcast.

clbroadcast is used as follows:

```
% clbroadcast [nodename ...]  
% clbroadcast -mp [nodename ...]  
% clbroadcast -telnet [nodename ...]
```

Examples

The following command broadcasts to cluster nodes "nodea", "nodeb", and "nodec" using the default (telnet) interface:

```
% clbroadcast nodea nodeb nodec
```

To broadcast to all cluster nodes using the MP interface:

```
% clbroadcast -mp
```

For more details on the usage of clbroadcast, invoke the command:

```
% man clbroadcast
```

ClusterPack



Service ControlManager (SCM) Overview

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

[3.2.1 What is ServiceControl Manager?](#)

[3.2.2 How to install, configure, manage, and troubleshoot SCM:](#)

[3.2.3 How to Run SCM Web-based GUI](#)

3.2.1 What is ServiceControl Manager?

ServiceControl Manager (SCM) makes system administration more effective, by distributing the effects of existing tools efficiently across nodes. The ServiceControl Managed Cluster establishes the domain across which these tools operate.

SCM Features:

- Role-based management allows delegation of specific management activities to specific individuals on specific nodes or node groups without having to provide root access to the individual.
- A Distributed Task Facility that improves operator efficiency by replicating operations across the nodes or node groups within the ServiceControl Managed Cluster with a single command.
- Tools designed to deal with a single system (single system aware tools), like bdf, are dispatched to the target systems and their results collected for review. This mechanism can also be used to handle custom tools such as user-developed scripts or applications. A number of tools encapsulating HP-UX commands are provided with ServiceControl Manager.
- Mechanisms for managing tools already designed to deal with multiple systems, such as Ignite-UX and Software Distributor. Such multiple system aware tools benefit from consistent node group-based targeting and role-based authorization across the tools.
- Node groups simplify distribution of management activities across collections of nodes. Node group names can be chosen based on applications, system configuration or other meaningful attributes.

[Back to Top](#)

3.2.2 How to install, configure, manage, and troubleshoot SCM:

ServiceControl Manager is installed as part of ClusterPack, and should not need to be installed manually.

For additional information about the configuration, management, or general troubleshooting, please refer to the ServiceControl Manager Technical Reference:

<http://docs.hp.com/hpux/onlinedocs/B8339-90030/B8339-90030.html>

[Back to Top](#)

3.2.3 How to Run SCM Web-based GUI

This release of ClusterPack includes a version of SCM that has a Web-based GUI.

To run the SCM GUI, point your Web browser at the following URL:

`% https://<manager_node_address>:50000/`

You must be using a recent version of Internet Explorer or Netscape in order to run the SCM GUI.

[Back to Top](#)

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

Copyright 1994-2004 hewlett-packard company

ClusterPack



System Inventory Manager (SIM) Overview

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

[3.3.1 What is System Inventory Manager?](#)

[3.3.2 How to invoke SIM](#)

3.3.1 What is System Inventory Manager?

The SIM application is a tool that allows you to easily collect, store and manage inventory and configuration information for the Compute Nodes in the HP-UX Itanium 2 cluster. The SIM GUI in conjunction with the underlying Desktop Management Interface (DMI) provides a easy web-based interface.

Some of the key features of the tool are:

- You design the grouping of devices in the way that best suits your environment.
- The GUI's buttons, tabs and menus provide quick access to:
 - defining devices, and groups
 - adding, configuring, and deleting devices as well as groups, schedules and filters
 - collecting data on a scheduled basis or on demand
 - filtering of collected data to isolate specific data
 - comparing collected inventory data for devices and snapshots
 - reviewing the wealth of information that is contained in collected snapshots in an easy-to-use format
 - viewing group, device, schedule, snapshot and filter properties
- The scheduling facility allows you to define when SIM retrieves inventory information, dynamically and on pre-determined schedules that suit your needs. Data can be collected for individual devices and for each defined group.

- The filtering facility allows you to define and view only the information that you need at any given time.
- The Command Line Interface (CLI) that is provided enables scripting capabilities.

Documentation for SIM is available at:

<http://software.hp.com/products/SIM/info.html>

Online help is also available by clicking the Help Tab in SIM GUI.

[Back to Top](#)

3.3.2 How to invoke SIM

Using the SCM GUI:

- Under "Tools", select "HP Systems Inventory Manager"
- Double-click on the "HP Systems Inventory Manager" icon
- This launches the SIM GUI

From your web browser at your desktop:

Go to:

`http://<management_server>:1190/simgui`

The user name is the name that will appear on the GUI.

[Back to Top](#)

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

ClusterPack



Application ReStart (AppRS) Overview

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

[3.4.1 What is AppRS?](#)

3.4.1 What is AppRS?

AppRS is a collection of software that works in conjunction with Platform Computing's Clusterware™ to provide a fail-over system that preserves the current working directory (CWD) contents of applications in the event of a fail-over. Many technical applications provide application-level checkpoint/restart facilities in which the application can save and restore its state from a file set. Checkpoint/restart is particularly helpful for long running applications because it can minimize lost computing time due to computer failure. The usefulness of this capability is diminished however by two factors. First, computer failure frequently leaves the restart files inaccessible. Using a shared file system does not preclude data loss and can introduce performance degradation. Redundant hardware solutions are often financially impractical for large clusters used in technical computing. Secondly, applications affected by computer failure generally require human detection and intervention in order to be restarted from restart files. Valuable compute time is often lost between the time that the job fails and a user is made aware of the failure. Clusterware™ + AppRS provides functionality to migrate and restart applications affected by an unreachable host and ensure that the content of the CWD of such applications is preserved across a migration.

AppRS is accessed by submitting jobs to AppRS-enabled queues. Such queues generally end in `_apprs`. A number of utilities are also available for monitoring a job and its files:

- `apprs_hist`
- `apprs_ls`
- `apprs_clean`
- `apprs_mpijob`

More information is available in the man page or HP Application ReStart User's Guide.

```
% man apprs
```

To use AppRS, users must add the following line to their `~/.cshrc` file:

```
source /share/platform/clusterware/conf/cshrc.lsf
```

and the following line to their `~/.profile` file:

```
. /share/platform/clusterware/conf/profile.lsf
```

References:

- [2.3.4 Remove temporary files from Compute Nodes](#)
- [2.3.5 Prepare application for checkpoint restart](#)
- [2.3.6 Restart application from a checkpoint if a Compute Node crashes](#)
- [AppRS Release Note](#)
- [AppRS User's Guide](#)

[Back to Top](#)

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

Copyright 1994-2004 hewlett-packard company

ClusterPack



Cluster Management Utility (CMU) Overview

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

[3.5.1 What is CMU?](#)

[3.5.2 Command line utilities](#)

[3.5.3 Nodes monitoring](#)

[3.5.4 Invoking CMU](#)

[3.5.5 Stopping CMU](#)

[3.5.6 CMU main window](#)

[3.5.7 Monitoring By Logical Group](#)

[3.5.8 Contextual Menu](#)

[3.5.9 Logical Group Administration Menu](#)

3.5.1 What is CMU?

CMU is designed to manage a large group of Compute Nodes. CMU comes with a Graphical User Interface. It provides access to all Compute Nodes from a single screen using a single mouse click. The CMU main window gives you access to all the menus you need to setup your CMU configuration.

[Back to Top](#)

3.5.2 Command line utilities

CMU offers several command-line-based procedures:

- Generic command line to start or stop CMU
- Automatic Terminal Server configuration procedure if needed
- CMU database backup and restore procedure

[Back to Top](#)

3.5.3 Nodes monitoring

- Cluster monitoring

Enhanced monitoring capabilities for up to 1024 nodes in a single window (with vertical scrollbars).

- Monitoring tools

Provides tools to monitor remote node activities.

- Node Administration

Allows execution of an action on several nodes with one command. The actions are:

1. Boot and reboot selected nodes.
2. Console broadcast on selected nodes (by telnet: through the management network, through its management card if any, or through its console port if there is a terminal server).
3. Connect to a node by telnet: through the management network, through its management card if any, or through its console port if there is a terminal server.

- Event handling management

Displays a warning message or executes a command when a node becomes unreachable or reachable again.

[Back to Top](#)

3.5.4 Invoking CMU

The user must be logged in as root and use the generic command line 'cmutool' in a terminal window to start the CMU GUI main window:

```
% /opt/cmu/scripts/cmutool start cmu
```

or

```
% /opt/clusterpack/bin/cmu
```

The main CMU window is displayed at your terminal with the monitoring by Logical Group

window enabled. CMU will display the last monitored logical group.

Note:

When starting the CMU window for the first time, the monitoring action is performed with the “Default” Logical Group.

Note:

Some of the menus and functions within CMU will allow the user to act on more than one selected item at a time. When appropriate, the user can select multiple items by using the Ctrl or Shift keys in conjunction with the left mouse button. The Ctrl key allows multiple non-contiguous selections while the Shift key allows contiguous or groups of objects to be selected.

[Back to Top](#)

3.5.5 Stopping CMU

To stop CMU, left click the mouse on the Quit button in the main CMU window (lower right corner).

Note:

When stopping, CMU saves the current configuration parameters.

[Back to Top](#)

3.5.6 CMU main window

Description of the main menu buttons for CMU monitoring and management:

- Monitoring Windows

Activates a menu to choose how the nodes are grouped. It contains the following options:

- Monitoring By Logical Group
- Monitoring By Terminal Server

- Configuration Windows

Activates a menu with the following options:

- Terminal Server Configuration
- PDU Configuration
- Network Topology Adaptation
- Node Management
- Event Handling Configuration

[Back to Top](#)

3.5.7 Monitoring By Logical Group

The following section describes the different actions that the user can perform in the "Monitoring By Logical Group" window.

- Select/Unselect one node

Left click on the name of this node. The node becomes darker when selected, or returns to original color when unselected.

- Select/Unselect several nodes

Left click on the name of the nodes one by one. If the node clicked is not selected, it will be added to your selection. If it is already selected, it will be removed from the selection. The selection is composed of all the darker nodes on the window.

- Select all the nodes of the logical group

Double-left click on one node of the logical group and all the nodes will be selected.

- Unselect all the nodes of the logical group

Double-middle click on one node of the logical group and all the nodes will be selected.

- Select a set of contiguous nodes of the logical group

Left click on the first node and hold the Shift key when selecting the last node of the group (with a left click). All the nodes between the two selected nodes will be added to the selection.

[Back to Top](#)

3.5.8 Contextual Menu

A contextual menu window appears with a right click on a node displayed in the central frame of the main monitoring CMU window. The following menu options are available:

- Telnet Connection

Launches a telnet session to this node. The telnet session is embedded in an Xterm window.

- Management Card Connection

Launches a telnet connection to the management card of this node. The telnet session is embedded in an Xterm window. This assumes that the management card property was properly set up during CMU installation. If the node does not have any management card, this menu is inaccessible.

- Locator On

Switches on the Locator LED of the node. This option is only available if the node is an HP Integrity server with a properly registered ECI card.

- Locator Off

Switches off the Locator LED of the node. This option is only available if the node is an HP Integrity server with a properly registered ECI card.

Note:

If several nodes are selected, all the items of the contextual menu are inactivated.

[Back to Top](#)

3.5.9 Logical Group Administration Menu

The Logical Group Administration item in the menu bar provides the following commands:

- Select/Unselect

The user can select/unselect one or several nodes by a left click in the node box in the central window frame. To select/unselect all the nodes in one step, use the corresponding entry in the "Logical Group Administration" menu.

Note:

Many management actions such as boot, reboot, halt, or monitoring will be applied to all of the selected nodes.

- Halt

This sub-menu allows a system administrator to issue the halt command on all of the selected nodes. The halt command can be performed immediately (this is the default), or delayed for a given time (between 1 to 60 minutes). The administrator can also have a message sent to all the users on the selected nodes by typing in the "Message" edit box.

Note:

The halt command is performed on the nodes using "rsh". This implies that on the Compute Node, permission must be given to perform commands as superuser/root from the Management Server. If not, the halt command will not work properly.

- Power Off

This feature allows the user to power off the nodes that have a management card, or that are linked to a remotely manageable PDU. All the nodes to be booted must have the same management card password and the same PDU password. If a node is linked with both a PDU and a management card, the power off will be performed using the management card. The PDU will be used only if the management card power off has failed.

Note:

If the nodes are not halted, they will be powered off by the remotely manageable PDU or by their management card. This can damage the file system. If unsure, use "Halt" before powering off.

- Boot

This feature allows the user to boot a collection of nodes on their own local disk or over the network. The user needs to select the node or nodes to be booted prior to launching this command. The boot procedure can use the management card of each node or the remotely manageable PDU connected to the nodes. In this case the password for the management card or for the PDU must be entered. All the nodes to be booted must have the same management card password and the same PDU password.

Note:

If the nodes are already up, they will be reset by the remotely manageable PDU or by their management card. This can damage the file system. If unsure, use "Halt"

before booting a node.

- Reboot

This sub-menu allows a system administrator to issue the reboot command on all of the selected nodes. The reboot command can be performed immediately (this is the default), or delayed for a given time (between 1 to 60 minutes). The administrator can also have a message sent to all the users on the selected nodes by typing in the "Message" edit box.

Note:

The reboot command is performed on the nodes using "rsh". This implies that on the Compute Node, permission must be given to perform commands as superuser/root from the Management Node. If not, the reboot command will not work properly.

- Console Broadcast

There are three types of connections for console broadcast:

- Telnet connection through the network, if the network is up on all selected nodes.
- Telnet connection to the console and through a terminal server, if all the selected nodes are connected to a terminal server.
- Telnet connection through the management card, if all the selected nodes have a management card.

Note:

Telnet connections through the management card are not allowed in a single window mode.

- Multiple Window

If the user chooses the multiple windows mode, the command launches a master console window and concurrent mirrored telnet sessions (embedded in an Xterm) on all the selected nodes. All input typed in the master console window is mirrored (broadcast) to the telnet sessions on all the selected nodes. This allows a system administrator to issue the same command on several nodes, while typing it only once.

In addition, all input typed into one of the individual telnet windows is directed only to the underlying node. This allows the issuing of commands specific to a node by typing input directly into the desired Xterm for that node.

To improve the Xterm windows display appearance, every window can be shifted (in x and y) from the previous one to make sure that they fit nicely on the screen. By default, the shift values are computed so that the windows tile the screen and no window is displayed outside of the screen.

If the user does not need to visualize the telnet sessions, or does not want to crowd the display, the user has the option to start the Xterm windows minimized.

Note:

The console broadcast displayed Xterm windows are limited by the number of ttys and the display capacity of the X server. HP advises the use of a Single Window for performing the broadcast command on a large number of nodes.

- Remote Connection

This feature offers the same service as a right click (in the CMU central frame) and allows the user to establish a telnet connection through the Ethernet network to a machine listed in CMU. If the connection fails, you must press a key to destroy the window.

[Back to Top](#)

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

Copyright 1994-2004 hewlett-packard company

ClusterPack



NAT/IPFilter Overview

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

[3.6.1 Introduction to NAT \(Network Address Translation\)](#)

3.6.1 Introduction to NAT (Network Address Translation)

Network Address Translation (NAT) or IP Aliasing provides a mechanism to configure multiple IP addresses in the cluster to present a single image view with a single external IP address.

The importance of this is two-fold:

(1) Single point of "controlled" access to the cluster Compute Nodes.

Network Address Translation allows communications from inside the cluster to get out, without allowing connections from outside to get in. NAT rewrites the IP headers of internal packets going out, making it appear that they all came from a single IP address (which is the external IP address of the entire cluster). Reply packets coming back are translated back, and forwarded to the appropriate Compute Node. Thus, the Compute Nodes are allowed to connect to the outside world, if needed. However, outside machines cannot initiate any connection to individual Compute Nodes, since they are exposed only to the "translated" IP address of the entire cluster.

(2) Simplified network administration.

The administrator can add or delete nodes to from the cluster compute cloud without any impact to the external world. Further, the administrator can have the entire compute cloud in a private IP sub-net (10.x.y.z range, 192.168.p.q range), and this also alleviates the problem of a shortage of IP addresses.

	IP Aliasing or Network Address Translation (NAT)
--	--

IP Aliasing or Network Address Translation (NAT)

ClusterPack comes with HP-UX IPFilter, a software component with powerful packet filtering and firewalling capabilities. One of the features that it supports is Network Address Translation.

For information on HP-UX HPFilter, please refer to the HP-UX HPFilter manual and release notes at docs.hp.com:

<http://docs.hp.com/hpux/internet/index.html#IPFilter/9000>

For information on NAT features of HP-UX HPFilter refer to the public domain how-to document. No guarantee can be made about the correctness, completeness or applicability of this or any third party information.

<http://www.obfuscation.org/ipf/>

[Back to Top](#)

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

Copyright 1994-2004 hewlett-packard company

ClusterPack



Platform Computing Clusterware Pro V5.1 Overview

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

[3.7.1 What is Clusterware Pro?](#)

[3.7.2 How do I obtain and install the Clusterware Pro V5.1 license file?](#)

[3.7.3 Where is Clusterware Pro V5.1 installed on the system?](#)

[3.7.4 How can I tell if Clusterware Pro V5.1 is running?](#)

[3.7.5 How do I start and stop the Clusterware Pro V5.1 daemons?](#)

[3.7.6 How do I start and stop the Clusterware Pro V5.1 Web GUI?](#)

[3.7.7 What system resources are required by Clusterware Pro V5.1?](#)

[3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?](#)

[3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?](#)

[3.7.10 How can the Clusterware Pro V5.1 services be refreshed after changes to the configuration are made?](#)

[3.7.11 Where can I find more information about using and administering Clusterware Pro V5.1?](#)

3.7.1 What is Clusterware Pro?

Platform Computing Clusterware Pro V5.1 is a comprehensive cluster management solution for enterprises looking to maximize on the cost-effective, high performance potential of HP-UX clusters. Platform Computing's Clusterware:

- Seamlessly incorporates industrial-strength distributed workload management, provisioning capabilities, and unified view of cluster-wide node status and resources into an easy-to-leverage platform
- Provides a single view of cluster health and performance
- Allows your organization to evolve along specific integration paths in order to provide a customized HP-UX cluster management environment
- Extends to meet the growing demands of your company with Grid-enabled scalability.

Platform Computing's Clusterware Pro V5.1 provides the following benefits

- Organizations experience increased productivity from transparent single system, cluster-as-server access to compute resources.
- The Platform Computing's Clusterware Pro V5.1 solution dramatically reduces time to market through continuous access to the cluster's compute power.
- The Platform Computing's Clusterware Pro V5.1 solution enables organizations to achieve higher quality results by running simulations and analyses faster than previously possible.
- The integrated resource management capability of the Platform Computing's Clusterware Pro V5.1 solution minimizes cluster downtime and ensures maximum availability of the compute resources used for critical design and research work.

[Back to Top](#)

3.7.2 How do I obtain and install the Clusterware Pro V5.1 license file?

Obtain a License File

If you have purchased ClusterPack Clusterware Edition, you will need a license for Platform Computing's Clusterware Pro. You can call, email or fax your request to Hewlett-Packard Software Licensing Service. Refer to your Software License Certificate for contact information.

You will need to get the host identification number from the Management Server. The host ID can be found using the uname command:

```
% /bin/uname -i
```

The number returned by the command uname -i must be preceded by a '#' when making your request. For example, if uname -i returns 2005771344, provide the ID number as #2005771344 in your key request.

Note:

It may take up to 24 hours to receive the license file, depending on how the license request is submitted. Please allow sufficient time to receive the license file before proceeding with the installation of the ClusterPack software.

Install the license file

After you have received the Clusterware Pro license file, it must be installed on the Management Server. This is normally done as part of manager_config. During manager_config, you will be asked to provide the location of the license file. The license file will be copied into the correct location.

If I have a DEMO license file

By default, Clusterware Pro is configured to use a permanent license key. If you have received a DEMO license file, you will need to install that file by hand. Detailed instructions are provided in:

```
/share/platform/clusterware/conf/README.demo
```

Setup and Configuration of a DEMO license

The use of a DEMO license file (license.dat) for Clusterware Pro, as part of the ClusterPack V2.3 Clusterware Edition, requires some modification of installed configuration files.

These modifications will have to be removed in order to use a purchased license key (LSF_license.oem).

1. Place the DEMO license key onto the Management Server

```
/share/platform/clusterware/conf/license.dat
```

2. Modify the /share/platform/clusterware/conf/lsf.conf file

- o Comment out the line for "CLUSTERWARE_LICENSE_FILE".
- o Add a line for "LSF_LICENSE_FILE"

```
LSF_LICENSE_FILE=/share/platform/clusterware/conf/license.dat
```

3. Touch the permanent license file

```
% touch /share/platform/clusterware/conf/LSF_license.oem
```

4. Start the Clusterware Services on the Management Server

```
% /share/platform/clusterware/lbin/cwmgr start
```

Note:

These changes will need to be undone in order to use a permanent license key. Please see /share/platform/clusterware/conf/README.demo for more information.

References:

- [Step 7 Configure the ProCurve Switch](#)
- [2.2.1 Transfer a file from intranet to the Management Server in the cluster](#)
- [3.7.2 How do I obtain and install the Clusterware Pro V5.1 license file?](#)

[Back to Top](#)

3.7.3 Where is Clusterware Pro V5.1 installed on the system?

Platform's Clusterware Pro V5.1 is installed on the Management Server and Compute Nodes in the following directory:

```
/share/platform/clusterware
```

This directory resides on the Management Server and is NFS mounted on all compute nodes.

The /etc/exports file on the Management Server, and the /etc/fstab file on each Compute Node is updated automatically by ClusterPack.

[Back to Top](#)

3.7.4 How can I tell if Clusterware Pro V5.1 is running?

On the Management Server, several Clusterware Pro V5.1 services must be running in order to provide functionality for the tool. All of these services are located in /share/platform/clusterware.

In order to check on the status of the services, the following command can be issued:

```
% ps -ef | grep clusterware
```

On the Management Server, typical output will look like this:

```
root 20156 20121 0 Aug 2 pts/0 0:00 /share/platform/clusterware/etc/gabd -2
root 20152 20110 0 Aug 2 ? 0:00 /share/platform/clusterware/1.0/hppa11-64/etc/pim
root 20116 1 0 Aug 2 ? 0:00 /share/platform/clusterware/1.0/hppa11-64/etc/sbatchd
root 20163 20116 0 Aug 2 ? 0:05 /share/platform/clusterware/1.0/hppa11-64/etc/mbatchd -d /share/bla
root 20110 1 0 Aug 2 ? 0:11 /share/platform/clusterware/1.0/hppa11-64/etc/lim
root 20113 1 0 Aug 2 ? 0:00 /share/platform/clusterware/1.0/hppa11-64/etc/res
```

On a Compute Node, Clusterware Pro V5.1 uses different services than on the Management node. The method of checking the currently running processes is the same:

```
% ps -ef | grep clusterware
```

On a Compute Node, typical output will look like this:

```
root 17426 1 0 Aug 2 ? 8:35 /share/platform/clusterware/1.0/hppa11-64/etc/lim
root 17427 1 0 Aug 2 ? 0:00 /share/platform/clusterware/1.0/hppa11-64/etc/sbatchd
root 17428 1 0 Aug 2 ? 0:00 /share/platform/clusterware/1.0/hppa11-64/etc/res
```

[Back to Top](#)

3.7.5 How do I start and stop the Clusterware Pro V5.1 daemons?

A set of wrapper scripts are provided by the ClusterPack solution to streamline the starting and stopping of services.

To START services on the Management Server

Issue the following command on the Management Server as the super user (i.e. root):

```
% /share/platform/clusterware/lbin/cwmgr start
```

To STOP services on the Management Server

Issue the following command on the Management Server as the super user (i.e. root):

```
% /share/platform/clusterware/lbin/cwmgr stop
```

To START services on ALL Compute Nodes

Issue the following command on the Management Server as the super user (i.e. root):

On the Management Server

```
% clsh /share/platform/clusterware/lbin/cwagent start
```

To STOP services on ALL Compute Nodes

Issue the following command on the Management Server as the super user (i.e. root):

On the Management Server

```
% clsh /share/platform/clusterware/lbin/cwagent stop
```

To START services on a single Compute Node

Issue the following command as the super user (i.e. root):

On the Management Server

```
% clsh -C <compute_node> /share/platform/clusterware/lbin/cwagent start
```

On the Compute Node

```
% /share/platform/clusterware/lbin/cwagent start
```

To STOP services on a single Compute Node

Issue the following command on the Compute Node as the super user (i.e. root):

On the Management Server

```
% clsh -C <compute_node> /share/platform/clusterware/lbin/cwagent stop
```

On the Compute Node

```
% /share/platform/clusterware/lbin/cwagent stop
```

References:

- [3.1.5 clsh - Runs commands on one, some, or all nodes in the cluster.](#)

[Back to Top](#)

3.7.6 How do I start and stop the Clusterware Pro V5.1 Web GUI?

The Web GUI is started and stopped as part of the tools that are used to start and stop the other Clusterware Pro V5.1 services. No additional steps are required.

Note:

The Clusterware Pro Web GUI is not automatically started during a reboot of the Management Server. The Web GUI must be manually started if the Management Server is rebooted.

References:

- [3.7.5 How do I start and stop the Clusterware Pro V5.1 daemons?](#)

[Back to Top](#)

3.7.7 What system resources are required by Clusterware Pro V5.1?

The Clusterware Pro V5.1 web server is Tomcat. Tomcat is maintained and distributed by the Free Software Foundation. Several tools within the ClusterPack solution use the Tomcat web server.

[Back to Top](#)

3.7.8 How do I access the Clusterware Pro V5.1 Web Interface?

The Clusterware Pro V5.1 Web Interface can be accessed from any web browser that is connected to the Management Server. The address is:

`http://<management_server_name>:8080/Platform/`

- Where `<management_server_name>` is the full domain name of the Management Server.

- The username and password are the same as for any normal user account on the Management Server.

References:

- [3.7.6 How do I start and stop the Clusterware Pro V5.1 Web GUI?](#)

[Back to Top](#)

3.7.9 How do I access the Clusterware Pro V5.1 Command Line Interface?

Before using the Clusterware Pro V5.1 CLI, you must set a number of environment variables. This must be done once in each shell before using any of the Clusterware Pro V5.1 commands.

- For users of csh or tcsh

```
% source /share/platform/clusterware/conf/cshrc.lsf
```

- For users of sh or bash or ksh

```
% . /share/platform/clusterware/conf/profile.lsf
```

[Back to Top](#)

3.7.10 How can the Clusterware Pro V5.1 services be refreshed after changes to the configuration are made?

The services only read the configuration file when they are started up or reconfigured. Any time a change is made to the configuration, the services must either be restarted or reconfigured. Changes include, but are limited to:

- adding or removing queues
- changing existing queues
- adding or removing nodes
- reinstalling software (i.e. manager_config -F)

Reconfiguring the Clusterware Pro V5.1 Services

The Clusterware Pro V5.1 services can be reconfigured without stopping the daemons. This must be done by the super user (i.e. root) and you must be able to run the 'rsh' command without providing a password. The 'f' flag will force a reconfiguration, unless there are fatal errors. When reconfiguring the cluster, order is important:

```
% lsadmin reconfig -f
```

```
% badmin reconfig  
% badmin mbdrestart -f
```

Restarting the Clusterware Pro V5.1 Services

As an alternative, the Clusterware Pro V5.1 services can simply be restarted on all nodes in the cluster. This will cause any information about jobs that are running to be lost, but the jobs will continue to run. Please "How do I start and stop the Clusterware Pro V5.1 daemons?" for more information.

References:

- [3.7.5 How do I start and stop the Clusterware Pro V5.1 daemons?](#)

[Back to Top](#)

3.7.11 Where can I find more information about using and administering Clusterware Pro V5.1?

Online reference documents are available for Administering Clusterware Pro, and Running Jobs using Clusterware Pro. These documents provide more detail on the commands that are part of the Online Tuto

[Administering Platform Clusterware Pro \(pdf\)](#)

[Running Jobs with Platform Clusterware Pro \(pdf\)](#)

[Back to Top](#)

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

Copyright 1994-2004 hewlett-packard company

Management Processor (MP) Card Interface Overview

ClusterPack



Management Processor (MP) Card Interface Overview

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

[3.8.1 Using the MP Card Interface](#)

3.8.1 Using the MP Card Interface

The MP cards allow the Compute Nodes to be remotely powered up. Using this technology, the initial installation and configuration of the Compute Nodes is eased.

In order to access the MP Card Interface (using HPUX 11i V2.0):

- Enter Ctrl-B from the system console (serial or LAN) to activate the main MP menu.
- Enter the 'cm' command to access the command menu.
 - Enter the 'pc' command (power control) to toggle system power state. Note that no signal is sent to the OS to allow for a graceful shutdown, so the system should be halted prior to using this command to turn off the system.
 - Enter the 'lc' command (LAN configuration) to set IP address, subnet mask, gateway address, and web console port number for the cluster LAN port.
 - Enter the 'xd -r' command (reset and diagnostics) command to reset the MP card.
 - Enter Ctrl-B from the system console (serial or LAN) to return to the main MP menu.
- Enter the 'x' command to completely disconnect from the LAN port.

References:

○

[Back to Top](#)

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

Copyright 1994-2004 hewlett-packard company

Related Documents

ClusterPack



Related Documents

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)

[4.1.1 HP-UX 11i Operating Environments](#)

[4.1.2 HP-UX ServiceControl Manager](#)

[4.1.3 HP Application ReStart](#)

[4.1.4 HP System Inventory Manager](#)

[4.1.5 HP-UX IPFilter](#)

[4.1.6 ClusterPack V2.3](#)

4.1.1 HP-UX 11i Operating Environments

HP-UX 11i March 2002 Release Notes

<http://www.docs.hp.com/hpux/onlinedocs/5185-4391/5185-4391.html>

<http://www.docs.hp.com/hpux/os/11i/index.html>

Ignite-UX Administration Guide

<http://www.docs.hp.com/hpux/onlinedocs/B2355-90772/B2355-90772.html>

Software Distributor Administration Guide for HP-UX 11i

<http://www.docs.hp.com/hpux/onlinedocs/B2355-90740/B2355-90740.html>

[Back to Top](#)

4.1.2 HP-UX ServiceControl Manager

HP-UX ServiceControl Manager 3.0 User's Guide

<http://www.docs.hp.com/hpux/onlinedocs/5187-4543/5187-4543.html>

ServiceControl Manager Troubleshooting Guide

<http://www.docs.hp.com/hpux/onlinedocs/5187-4198/5187-4198.html>

[Back to Top](#)

4.1.3 HP Application ReStart

HP Application ReStart Release Note

[AppRS Release Notes \(pdf\)](#)

HP Application Restart User's Guide

[AppRS User's Guide \(pdf\)](#)

[Back to Top](#)

4.1.4 HP System Inventory Manager

SIM Info

<http://software.hp.com/products/SIM/info.html>

[Back to Top](#)

4.1.5 HP-UX IPFilter

HP-UX IPfilter Release Note

<http://www.docs.hp.com/hpux/onlinedocs/B9901-90010/B9901-90010.html>

[Back to Top](#)

4.1.6 ClusterPack V2.3

ClusterPack



Dictionary of Cluster Terms

[Cluster LAN/Switch](#)

[Cluster Management Software](#)

[Guarded Cluster](#)

[Head Node](#)

[Interconnect Switch](#)

[Management Processor \(MP\)](#)

[Management Server](#)

[Network Attached Storage \(NAS\)](#)

[Storage](#)

Cluster LAN/Switch

A Cluster LAN/Switch is usually an Ethernet network used to monitor and control all the major system components. May also handle traffic to the file server.

The Cluster Management Software is the ClusterPack for system administrators and end-users.

[Back to Top](#)

Guarded Cluster

A cluster where only the Management Server has a network connection to nodes outside of the cluster. All of the Compute Nodes are connected within the cluster on a private subnet (i.e. IP addresses of 10.*.*.* or 198.162.*.*).

[Back to Top](#)

Head Node

A Head Node provides user access to the cluster. In smaller clusters, the Management Server may also serve as a Head Node.

References:

- [Management Server](#)

[Back to Top](#)

Interconnect Switch

An Interconnect Switch provides high speed connectivity between Compute Nodes. Used for message passing and remote memory access capabilities for parallel applications.

[Back to Top](#)

Management Processor (MP)

Management Processor (MP) controls the system console, reset and power management functions.

[Back to Top](#)

Management Server

The Management Server provides single point of management for all system components in the cluster. In smaller clusters the Management Server may also serve as a head Node.

References:

- [Head Node](#)

[Back to Top](#)

Network Attached Storage (NAS)

Network Attached Storage (NAS) attaches directly to Ethernet networks, providing easy installation, low maintenance, and high uptime.

[Back to Top](#)

Storage

Storage can either be local to each Compute Node, or external to the cluster.

[Back to Top](#)

[Index](#) | [Administrators Guide](#) | [Users Guide](#) | [Tool Overview](#) | [Related Documents](#) | [Dictionary](#)
